# Concept pluralism in conceptual engineering

Sarah Sawyer

Published online: 04 Oct 2021.

Submit your article to this journal ☑

View related articles ☑

View Crossmark data ☑

# Concept pluralism in conceptual engineering

Sarah Sawyer

Philosophy, University of Sussex, Brighton, UK

**ABSTRACT**
In this paper, I argue that an adequate meta-semantic framework capable of accommodating the range of projects currently identified as projects in conceptual engineering must be sensitive to the fact that concepts (and hence projects relating to them) fall into distinct kinds. Concepts can vary, I will argue, with respect to their direction of determination, their modal range, and their temporal range. Acknowledging such variations yields a preliminary taxonomy of concepts and generates a meta-semantic framework that allows us both to accommodate the full range of cases and to identify a proper subset of concepts for special ameliorative consideration. Ignoring such variations, in contrast, leads to a restricted meta-semantic framework that accommodates only a subset of the particular projects while generating implausible accounts of others.

## 1. Introduction

There are two distinguishable strands to the literature on conceptual engineering. The first reflects the development of particular projects related to phenomena such as truth, gender and race. The second reflects the development of alternative meta-semantic frameworks within which those particular projects might best be understood. The two strands are clearly related: specific projects in conceptual engineering typically make meta-semantic assumptions; and meta-semantic frameworks often appeal to specific examples for support. This paper is a contribution to the second strand; but it is also an attempt to address the more general question of the relation between the two. I will argue that an adequate meta-semantic

**CONTACT** Sarah Sawyer ✉ s.a.sawyer@sussex.ac.uk 🖅 Philosophy, University of Sussex, Brighton BN1 9RH, UK

framework capable of accommodating the range of projects currently identified as projects in conceptual engineering must be sensitive to the fact that concepts (and hence projects relating to them) fall into distinct kinds. Concepts can vary, I will argue, with respect to their direction of determination, their modal range, and their temporal range. Acknowledging such variations yields a preliminary taxonomy of concepts and generates a meta-semantic framework that allows us both to accommodate the full range of cases and to identify a proper subset of concepts for special ameliorative consideration. Ignoring such variations, in contrast, leads to a restricted meta-semantic framework that accommodates only a subset of the particular projects while generating implausible accounts of others.

The structure of the paper is as follows. In section 2, I provide an overview of a meta-semantic framework that I've developed and defended elsewhere.[1] It's an externalist framework that involves the three key notions of meanings, concepts and conceptions, and is important for what I'll go on to say later in the paper. In section 3, I explain what I mean by a concept's direction of determination; in section 4, I explain what I mean by a concept's modal range; and in section 5, I explain what I mean by a concept's temporal range. In section 6, I discuss the limitations of alternative meta-semantic frameworks with respect to accommodating the full range of cases. I conclude briefly in section 7.

## 2. Meanings, concepts and conceptions

I start by rejecting a common assumption, namely the assumption that the semantic content of a sentence is the same as the propositional content of the thought it expresses. For example, the semantic content of the sentence 'Tigers have stripes' is typically understood as the proposition that tigers have stripes, which proposition is also typically taken to provide the (propositional) content of the thought the sentence expresses. So semantic contents and thought contents are typically identified.[2] The identification is also evident in indexical and demonstrative cases, since the content of a subject's utterance of 'I am here now', for example, is taken to be identical to the content of the thought she expresses, even though that content might vary from person to person, place to place, and time to time.

---

[1]See for example Sawyer (2018, 2020a, 2020b).
[2]The distinction between language and thought is drawn in Burge (1986); and the distinction between linguistic meaning and Fregean sense is articulated in Burge (1979b, 1990). My view can be seen as a development of Burge's, and hence as Fregean in spirit.

At the level of words rather than sentences, the common assumption is that the term 'meaning' and the term 'concept' are harmlessly interchangeable (perhaps even synonymous). This is evident in the widespread practice amongst philosophers of switching freely between the two, often within a single sentence, as if the difference were merely a matter of terminological preference. Those who do not switch back and forth tend to talk exclusively in terms of meanings, eschewing talk of concepts altogether, sometimes on the grounds that concepts are metaphysically obscure, but, ultimately, on the grounds that talk of concepts is otiose. Cappelen's *Fixing Language* provides an explicit example of this practice from within the literature on conceptual engineering, but the approach is evident in semantic accounts of conceptual engineering more generally.[3] The focus on meanings to the exclusion of concepts is, of course, a legitimate strand of the broader philosophical literature in meta-semantics, but it would be a mistake to think that meta-semantic accounts reveal the nature of thought.[4] In contrast to all such approaches, I think that meanings and concepts are ontologically distinct, and that a single term or expression will be associated with one of each; it will have a meaning, and it will also express a concept. I talk here of a term's *having* a linguistic meaning and *expressing* a concept, not because there is a difference between the *having* relation and the *expressing* relation, but merely to emphasise the difference in the relevant *relata* —in what is had, and in what is expressed, respectively.

One way to understand meanings and concepts is as distinct representational elements. This is the way I have put it in previous work. An alternative would be to reserve the expression 'representational element' for concepts—the fundamental representational constituents of thoughts—and to characterise meanings not as representational elements *per se,* but as encodings of linguistic practices. There are distinct pragmatic advantages to expressing matters in the second of these ways. However, they are broadly equivalent and the difference has no substantive impact on the main thread of the current paper. For now, I outline my account of meanings, and then my account of concepts.

---

[3]See Cappelen (2018). See also, for example, Ludlow (2014), Pinder (2020), and Plunkett and Sundell (2020).

[4]This kind of error is evident most recently in Yli-Vakkuri and Hawthorne (2018), for a response to which see Sawyer (2020c).

## 2.1. Meanings

I take the meaning of a term to be determined by a combination of patterns of actual use across a linguistic community together with relations of deference (typically implicit) between its members. The notion of deference serves to hold together the variable and often idiosyncratic uses of a single term by different members of a single practice. The notion of deference is thus designed to pick up the slack between actual individual use, which can vary dramatically from person to person, and public meaning, which is (relatively) stable across the community.[5] A reasonable notion of deference must of course recognise that the use of a term by some members of the community carries more weight than the use of that term by others—that not all uses contribute equally to the determination of meaning. But, despite the popularity of the phrase, it is best not to construe the role of deference in determining meaning as 'deference to experts'. This is for at least three reasons.

First, talk of experts conflates two distinct ideas: one is that of an expert as someone who is in the best epistemic position relative to other epistemic agents; the other is that of an expert as someone who is in the best epistemic position relative to the facts—that is, as someone who actually knows. It is the first of these that captures our ordinary use of the term 'expert'; but being an expert in this first, ordinary sense is clearly no guarantee that one is an expert in the second, fact-knowing sense. And yet taking experts to be experts in the second, fact-knowing sense has been the norm in the literature on semantic externalism (understood as a thesis about linguistic meaning). This is to some extent due to the examples that were used to introduce the notion of linguistic deference. For example, Putnam claims to mean *elm* by 'elm' and *beech* by 'beech' because although he cannot distinguish elms from beeches, he can defer to others who can. The people to whom he defers are taken to be knowers. But more generally, the Twin Earth thought experiments designed to promote semantic externalism themselves depend on the presence of a presumed knower, either as a protagonist, such as Alf's doctor in Burge's arthritis example, or as a theorist, such as we (and the scientists) are taken to be when assessing Putnam's original water

---

[5]The distinction here between an individual's use of a term and its public meaning is not the same as the distinction between speaker meaning and semantic meaning. Space precludes further elaboration of this point, but I suspect it could be used to counter the claim, found in Pinder (2020), that conceptual engineering can (and should) be understood in terms of speaker meaning.

example.[6] It is ironic, from my perspective, that the role of knowers has been such a feature of the debate given that content externalism depends not on the actuality of knowledge but on the possibility of ignorance. I return to this point later.

Second, talk of experts is apt to mislead one into thinking that patterns of deference flow in a single, clearly-defined direction, from the non-experts to the experts. But linguistic deference is more fluid and less discrete than this simple picture recognises. This is because the status of a speaker as someone to whom others defer is contingent upon the dispositions of those others, where no speaker is immune from correction, with the result that their status as one to whom others defer is subject to change in the light of new discoveries and insights by others. This is precisely because those to whom we defer in any given area need not be experts in the fact-knowing sense.

Third, talk of experts is liable to create the false impression that deference is restricted to a domain of specialist (for example, chemical, biological, or medical) terms, where the notion of an expert makes some literal sense. But deference plays a role in determining the meaning of every term in the language, from 'elm' and 'arthritis' to 'hipster' and 'banter'. This can be seen in its starkest form by reflecting on the way in which an infant learns to speak. The process is gradual; it depends on the 'teacher' reinforcing what they take to be the correct use of a word and correcting what they take to be the incorrect use of a word, where the 'teacher' need not themselves be an expert in either of the two senses described above. But this kind of process takes place for every word in the language, which illustrates the fact that deference has a maximally broad range. For now, I leave the notion of deference at an intuitive level, noting merely that it should not be construed in terms of deference to experts.[7]

For present purposes my interest lies exclusively with predicative terms, where the meaning of a predicative term, precisely because it is determined by patterns of actual use across a community and relations

---

[6]See Putnam (1970) for the elm/beech example as well as the water/twin water example. See Burge (1979a) for the arthritis example.

[7]In previous work I have appealed to the notion of 'deference to the most competent', echoing Burge's use of the phrase in his (1986). But in Burge's account 'the most competent' are those to whom others would defer under idealized dialectical conditions, which I now think too far removed from actual use to underpin an account of linguistic meaning. Actual relations of deference are subject to power-relations, and recognized indicators of competence need not be reliable indicators of actual competence. These facts can nonetheless influence the actual meanings of our terms. The account therefore needs to be modified, particularly, as will become clear later, when a concept's temporal range is taken into account.

of deference between its members, reflects the community's current, collective, received understanding of the relevant subject matter. By 'subject matter', I mean something intuitive and relatively weak, so that the subject matter of the term 'tiger' is tigers, and the subject matter of the term 'hipster' is hipsters. 'Subject matter' is not a technical expression; it is just intended to mark out whatever it is that the term allows us to talk about. Thinking of meaning as reflecting a community's current, collective, received understanding of a subject matter underpins the fact that meaning is determined in the context of relational, contextual applications to things in the world, whether that's to tigers or to hipsters.

The account of meaning just outlined is externalist in the sense that an individual's actual use of a term fails to determine its meaning. I end my discussion by emphasising four key elements of the account. First, the linguistic meaning of a term can be stable across a community of people whose individual uses differ. Second, actual as well as merely possible variation in individual use of even the most basic terms in a community is widespread and should not be set aside as atypical, as restricted to specialist terms, or as evident only in the very young or the cognitively challenged. The variation in our use of terms is underpinned by the variation in our perspectives on the world, where that variation is itself endemic and widespread. Third, linguistic meaning is messy: linguistic meaning can change over time, meanings can be contested, there can be no settled meaning for a term at a given time, and so on, and each of these maps onto the messiness of linguistic practice. Fourth, since the meaning of a term reflects the community's current, collective, received understanding of the relevant subject matter, meaning ultimately rests on a kind of *de facto* implicit agreement about the nature of the world at the time.[8]

## 2.2. Concepts

I take the concept expressed by a term to be individuated, fundamentally, by relations between a thinker and objective properties with instances of which she stands in an appropriate causal chain. Paradigm cases are provided by the classic Twin Earth thought experiments involving concepts such as *gold*, *water*, and *tiger*.[9] The thought experiments illustrate the fact that had our world been different in certain respects, we would

---

[8]This is not to be understood as a minimal moment in time, but rather as a period of time sufficient to accommodate the fact that utterances and reasoning take place in time.
[9]For which see Putnam (1970) and Kripke (1972).

have had different concepts, independently of whether we either had at the time, or acquired at a later time, sufficient epistemic capacities to distinguish between the kinds to which our concepts actually refer and the twin kinds to which our counterfactual twins' concepts refer. This is what makes the account an externalist account: the concepts a subject possesses are individuated neither by her individual epistemic capacities nor by her individual conceptions (her sets of associated beliefs). Oscar and Twin Oscar have different concepts—*water* and *twater* respectively—despite having the same epistemic capacities and associated conceptions narrowly construed;[10] in contrast, Oscar in 1750 and the scientists later on are assumed to have the very same concept—*water*—despite having different epistemic capacities and associated conceptions.

We saw above that the linguistic meaning of a term can be stable across a community of people whose individual uses differ. For parallel reasons, the concept expressed by a term can be stable across a community of people whose individual associated conceptions differ. But concepts have an additional level of stability that meanings do not have. The concept expressed by a term can be stable not only across variations in associated conceptions at the individual level, but also (for at least some concepts) across variations in associated conceptions at the communal level. This is because a concept is determined, fundamentally, by neither. Moreover, since the communal conception associated with a concept can be understood as the set of beliefs that constitute the community's current, received understanding of the relevant subject matter, the communal conception associated with a concept will map directly onto the linguistic meaning of the term by means of which the relevant concept is expressed. Thus, the concept expressed by a term can be stable across variations in linguistic practice at the level of the community, and hence across variations in linguistic meaning itself.[11]

For illustrative purposes, suppose that Community A and Community B (whether in different places or at different times) use term T to talk about a given subject matter, but that they have different associated communal conceptions of the subject matter in question. In such a

---

[10]Talk of epistemic capacities and associated conceptions 'narrowly construed' is not to be taken literally, since both epistemic capacities and associated conceptions involve representational states and hence cannot, I think, be construed narrowly (cf. Sawyer 2007). This kind of talk is merely intended to convey the kind of similarity between us and our counterfactual twins that Putnam aimed to identify.

[11]Temporal externalists maintain that linguistic meaning can remain stable across variations in linguistic practice (cf. Ball 2020; Jackman 1995, 2005). In this respect, temporal externalist accounts of linguistic meaning might be thought comparable to my account of concepts. However, since such accounts see linguistic meaning as determined by future linguistic practice, they nonetheless fail to provide the kind of practice-transcendent stability that I maintain some concepts exhibit (cf. Sawyer 2020d).

scenario, Community A and Community B, having different associated communal conceptions of the relevant subject matter, will also mean different things by their respective uses of term T, precisely because of the relation, noted above, between a linguistic practice and the communal conceptions that shape it. Nonetheless, this does not imply that Community A and Community B express different concepts by term T. Indeed, if the description of the communities as having different associated communal conceptions of the same subject matter is correct, there is reason to think the very same concept will be expressed by each community despite the differences in associated conceptions. Just as individual conceptions can be wrong (we can go wrong individually), communal conceptions can be wrong (we can go wrong collectively). That is, since the agreed characterisation of a subject matter need not be a true characterisation of the subject matter, the linguistic meaning of a term cannot, as a matter of unrestricted principle, be identified with the concept it expresses. The correct explication of a concept would provide a true characterisation of the subject matter, but the correct explication of a concept, and hence a true characterisation of the subject matter, need not be something upon which a community has yet alighted; nor need it be something upon which a community ever alights.

The distinction between meanings and concepts entails a corresponding distinction between linguistic norms and conceptual norms. Not only is there a way in which we ought to use a term given the linguistic practice of our community, but there is also a way in which we ought to use a term given the concept it expresses. These can diverge. Linguistic norms are generated by agreement and provide a standard against which individual use can be evaluated relative to the wider linguistic practice. Such practice-relative evaluation helps to promote uniformity in use and hence helps to promote communication. But conceptual norms are generated by truth and provide an absolute standard against which to evaluate not only individual use but also communal use. Thus an utterance of the sentence 'Whales are not fish' in 1700 would have been subject to legitimate practice-relative correction because it would have contravened the linguistic norms of the time; but, as we now recognise, the sentence expressed a true thought nonetheless. When people talk of truth being the norm of assertion, it is absolute conceptual norms rather than practice-relative linguistic norms that are in play; what we ought to assert is sentences that express true thoughts, not merely sentences that accord with linguistic practice.

## 2.3. Conceptions

I will finish this section by saying something about the explanatory role of individual conceptions in my framework. The role of concepts is, as should be clear, to secure objective reference; concepts, at the fundamental level, provide representational anchors to the world in a broadly theory-transcendent way.[12] This means concepts can provide stability through a change in linguistic practice and hence through a change in linguistic meaning. For example, a change in the linguistic practice surrounding the use of the term 'fish' would, in certain circumstances, constitute a change in linguistic meaning, since meaning depends on use; but that's consistent with the term's expressing the single concept—*fish*—both before and after the relevant change in linguistic practice. The role of individual conceptions, in contrast, is to explain and (in one sense) rationalise individual behaviour.[13] A child's individual conception of a fish, for example, can explain why she says 'Look at that gigantic fish' when she sees a whale.

Concepts and conceptions thus face in different directions: concepts face outwards, to the world, securing objective reference; conceptions face inwards, reflecting the relations between a subject's beliefs and other mental states. Understanding a subject's conceptions hence provides a partial insight into why she acts as she does; it allows us to see her actions as (again, in one sense) rational. Content internalist theories rightly emphasise the need to explain and rationalise individual behaviour in this way, and they rightly maintain that externally-individuated concepts are inadequate to the task—that reference to a subject's individual perspective is required instead. But they are wrong to think that concepts themselves need be individualistically-individuated as a result. The internalist's mistake is to conflate concepts with conceptions. And it is this conflation that undermines the very possibility of objective reference in the internalist framework.[14] Despite the fact that explanations of a subject's behaviour must, at times, make reference to her individual conceptions, such individual conceptions must themselves be composed of externally-individuated concepts if they are to be representational. It is concepts, rather than conceptions, that are

---

[12]As we will see, some concepts do this in a more direct way than others, where this depends in part on their direction of determination.

[13]The relevant sense here is rationalisation from the individual's perspective, but this does not imply a commitment to reasons internalism (cf. Sawyer 2014).

[14]See Sawyer (2007). See also Burge (2014).

the fundamental representational elements. But a theory of mind requires both.[15]

## 3. Directions of determination

I claimed at the outset that concepts can vary with respect to a number of properties. Specifically, they vary with respect to their direction of determination, their modal range, and their temporal range. In this section I explain what I mean by a concept's direction of determination, which, for any given concept, is either 'world-to-mind' or 'mind-to world'.

Some concepts have a 'world-to-mind' direction of determination. Concepts such as *gold* and *tiger* fall into this category. They depend for their individuation on the nature of the objective properties to which they refer. This is not to say that possession of a concept with a world-to-mind direction of determination does not depend for its individuation in some way on the cognitive states of individuals; it does. The point is rather that such concepts are not fully individuated by our conceptions, either individual or communal. We could all be wrong about the nature of gold and yet still possess the concept *gold*; we could all be wrong about the nature of tigers and yet still possess the concept *tiger*. A concept's having a world-to-mind direction of determination depends, at the most fundamental level, on the possibility of communal error concerning its referent, since interactions (of the right kind) with instances of the referent itself must be capable of circumventing both incorrect and incomplete conceptions of it. Such interactions with the referent contribute to determining that the concept possessed is the concept in question.

Determination conditions are not to be confused with acquisition conditions. An individual may acquire a concept with a world-to-mind direction of determination not by interacting with instances of the property to which the concept refers, but via a communicative act involving a characterisation of the relevant property. Since the characterisation provided may be either incorrect or insufficiently specific to be a uniquely identifying description of the relevant property, and the individual acquire the concept nonetheless, it is clear that the concept is acquired via the

---

characterisation but not determined by it. Concept acquisition is essentially an individual matter; concept determination, in contrast, is not essentially an individual matter. A single concept can be had by different members of a community each of whom has acquired the concept in a historically distinct way.

Other concepts have a 'mind-to-world' direction of determination. Concepts such as *game* and *juice* fall into this category: they ultimately depend for their individuation on our associated communal conceptions. It's not that games and juice aren't real things; it's just that what counts as a game or as juice depends on how we (as a community) use the relevant terms. This is what rules out the possibility of mass communal error. If community A's practice dictates that the term 'juice' is to be applied to fruit juice alone, and community B's practice dictates that the term 'juice' is to be applied to any kind of soft drink, including not only fruit juice and vegetable juice but also squash and lemonade, then the communities not only mean different things by the term 'juice', since their linguistic practices differ, but they also express different concepts by the term 'juice'. There is no independent, objective property that either community is trying to characterise and that could serve as an anchor for a single concept around which the linguistic practices of the different communities might vary. This means that the concepts expressed by such terms coincide with their meanings in the sense that the extensions of the relevant concepts are determined by the associated linguistic practices.[16] For such concepts, there are practice-relative norms against which individual use can be corrected, but there are no absolute norms over and above the practice-relative ones. For such concepts, adjudication between communities on grounds of truth makes no sense, although pragmatic considerations may, of course, have a bearing.

We can define what it is for a concept to have a world-to-mind direction of determination and what it is for a concept to have a mind-to-world direction of determination as follows:

> (WTM) A concept with a world-to-mind direction of determination is one which is not determined ultimately by our conceptions (individual or communal), and hence refers to a property about which the non-indexical beliefs of every member of the linguistic community could be false or incomplete.

---

[16]One way to put this is to say that the extensions of the concept and the meaning are necessarily identical. The concept and the meaning themselves will not be identical, since in general meanings are descriptive whereas concepts are not; but they are not distinct in the significant way in which a concept with a world-to-mind direction of determination is distinct from the linguistic meaning of the term by means of which that concept is expressed.

(MTW) A concept with a mind-to-world direction of determination is one which is determined ultimately by our conceptions (individual or communal), and hence refers to a property about which at least some of the non-indexical beliefs of at least some members of the linguistic community must be true.

Each principle consists of a metaphysical component and an epistemic corollary.[17] The epistemic corollaries are specified in terms of *non-indexical* beliefs because *indexical* beliefs provide a way for individuals (and communities) to identify instances of a property in advance of characterising it. Although it is plausible to think that individuals (and communities) could err in their identification of some instances of a property, it is implausible to think an individual (or a community) could err in their identification of every instance of a property and yet nonetheless possess a concept that refers to the property in question. This means that at least some of the relevant *indexical* beliefs of at least some members of the linguistic community must be true no matter what direction of determination the relevant concept has. The distinction between (WTM) and (MTW), then, concerns *non-indexical* beliefs, not indexical ones.

The principles, as stated, are mutually exclusive and jointly exhaustive, but they are silent on the actual classification of specific concepts, and there will undoubtedly be disagreement over cases, with realists about a given property referred to by a given concept maintaining that the relevant concept has a world-to-mind direction of determination, and anti-realists about that property maintaining the opposite. Disagreements along realist/anti-realist lines are disagreements about the nature of the property to which a concept is taken to refer, but are also, at the same time, disagreements about the nature of the concept itself. There will be some who disagree with even my basic examples, for example taking natural kind terms to have a mind-to-world direction of determination. But to think that there could be concepts which are neither themselves concepts with a world-to-mind direction of determination, nor constitutively related to concepts with a world-to-mind direction of determination, signifies a retreat to internalism. An externalist understanding of concepts, in contrast, implies that concepts with a mind-to-world direction of determination must be related constitutively (either directly or indirectly) to concepts with a world-to-mind direction

---

[17]The different epistemic relations we bear to natural kinds on the one hand and social kinds on the other is discussed in detail in Thomasson (2003). The view articulated resonates nicely with key aspects of this paper.

of determination.[18] Externalism is thus inconsistent with general forms of anti-realism.

The constitutive connection between concepts with a mind-to-world direction of determination and concepts with a world-to-mind direction of determination lies behind the qualification in (MTW) that a concept with a mind-to-world direction of determination is one which is determined *ultimately* by our conceptions. The qualification is intended to capture the idea that two communities cognitively identical with respect to their world-to-mind concepts could nonetheless differ with respect to their mind-to-world concepts as a result of relevant differences between their linguistic practices. It is not intended to convey the falsehood that world-to-mind concepts are not involved in the individuation of mind-to-world concepts. The non-cognitive world helps to individuate all of our concepts, but it does not ultimately determine them all.

## 4. Modal range

Just as concepts can differ with respect to their direction of determination, they can also differ with respect to their modal range. Concepts fall into one of three kinds in this respect; concepts can be 'practice-bound', 'world-bound', or 'world-invariant'. Any concept with a mind-to-world direction of determination is practice-bound because it is determined, ultimately, by our conceptions, and hence by our practice, whether that be individual or communal, linguistic or non-linguistic. I will say more about practice-bound concepts in the next section when I explain what I mean by a concept's temporal range, but for now I leave practice-bound concepts to one side. In this section I focus primarily on the distinction between world-bound concepts and world-invariant concepts, where this distinction marks a division amongst concepts with a world-to-mind direction of determination.

Concepts such as *gold, water,* and *tiger* are world-bound concepts. It is a necessary condition on a concept's being a world-bound concept that it can be possessed by an individual even though it refers to a property that has (actual or merely possible) epistemic counterparts for that individual at the time. This potential for epistemic counterpart referents is what explains the fact that world-bound concepts are subject to the original

---

[18]This is clearly true for certain explicitly composite concepts, such as *the possible person next door*, which contains the world-to-mind concepts *person,* and *possible*; but the claim that mind-to-world concepts are related constitutively to world-to-mind concepts is intended to capture the general truth that thought in general depends on the possession of world-to-mind concepts.

kind of Twin Earth thought experiment offered by Putnam; water has a potential epistemic counterpart in twin water, for example. What is relevant here is not the possibility of epistemic counterparts of instances of the property to which a concept refers; it is the deeper possibility of epistemic counterparts of the property itself. Thus gin is *not* an epistemic counterpart of water, despite the fact that it is possible to mistake a glass of gin for a glass of water on a particular occasion. And the claim is not that only world-bound concepts refer to properties with potential epistemic counterparts. Some practice-bound concepts do too, although there is, perhaps, a difference in scope. The referent of a world-bound concept is a property that has potential epistemic counterparts for every individual who possesses the relevant concept at the time. This connects with the fact that a world-bound concept refers to a property about which the non-indexical beliefs of every member of the linguistic community could be false or incomplete. In contrast, since a mind-to-world concept refers to a property about which at least some of the non-indexical beliefs of at least some members of the linguistic community must be true, it may be that the referent of a practice-bound concept must be a property that has potential epistemic counterparts for at most some individuals who possess the relevant concept at the time. The point is merely speculative, however, and does not affect the main line of argument.

The significant contrast is with world-invariant concepts, which refer to properties that do not admit of epistemic counterparts for anyone. I take mathematical concepts such as *number* and *successor*, epistemic concepts such as *knowledge* and *warrant*, normative concepts such as *moral goodness* and *justice*, and metaphysical concepts such as *causation* and *identity,* to be world-invariant concepts. We can be wrong in our identification of particular instances of these properties, and hence the particular instances have potential epistemic counterparts, but there is no deeper possibility of epistemic counterparts to the properties themselves. Because world-invariant concepts refer to properties that do not admit of epistemic counterparts, they are not subject to Twin Earth thought experiments, at least not of the same kind as the ones to which world-bound concepts are subject. The Moral Twin Earth thought experiment, as set out in Horgan and Timmons (1991), might be thought to tell against my claim here. But on a little reflection, it is clear that the Moral Twin Earth thought experiment is in fact an argument against the possibility of a Moral Twin Earth; it is an argument to the effect that moral terms such as 'good' and 'right' do not function like natural kind terms such as 'gold' and 'water'. I agree. The reason for this, however, is not

that an externalist semantics cannot be made to work for moral terms, as Horgan and Timmons suggest; rather, it is that the externally-individuated concepts expressed by moral terms such as 'good' and 'right' are world-invariant concepts, whereas the externally-individuated concepts expressed by natural kind terms such as 'gold' and 'water' are world-bound concepts.

World-invariant concepts should not be understood as referring to properties that are instantiated in every possible world; this is too strong. Some world-invariant concepts, such as *identity,* do plausibly refer to properties that are instantiated in every possible world. But other world-invariant concepts clearly do not. *Knowledge* and *warrant* refer to properties that depend on cognitive states, which some worlds lack. *Justice* and *moral goodness* refer to properties that depend not only on cognitive states but on particular kinds of cognitive creatures, which some worlds lack. There may be worlds in which there are no causal relations, and there may be worlds in which there are no asymmetric supervenience relations. But all of the relevant concepts are world-invariant nonetheless. Nor is there a commitment to the claim that world-invariant concepts can be acquired without interaction with instances of the property to which they refer; such a claim is implausible if taken as a universal principle. There is reason to think we acquire basic moral concepts by causal encounters with actions that instantiate moral properties, and there is reason to think we acquire basic mathematical concepts by causal encounters with groups of objects that instantiate mathematical properties. Rather, the distinction between world-bound concepts and world-invariant concepts lies solely in the different ways in which the possibility of communal error concerning their referents can arise. Communal error does not depend on potential epistemic counterparts.

We can define what it is for a concept to be world-bound and what it is for a concept to be world-invariant as follows:

> (WBC) A world-bound concept is one which has a world-to-mind direction of determination and refers to a property that has potential epistemic counterparts for individuals who possess it.

> (WIC) A world-invariant concept is one which has a world-to-mind direction of determination and refers to a property that does not have epistemic counterparts for individuals who possess it.

The principles, once again, reflect an interplay between metaphysical and epistemic claims. And again, as stated, they are mutually exclusive

and jointly exhaustive but silent on the actual classification of specific concepts. Here too, then, there may be disagreement over cases, and here too the disagreements will relate to disagreements between certain kinds of realists and anti-realists.[19] It might be tempting to think that the category of world-bound concepts is extensionally equivalent to the category of natural kind concepts, but this is not obviously the case. Concepts such as *belief* and *representation* are plausibly natural kind concepts but plausibly world-invariant nonetheless. This may be because they are also normative concepts, it may be because they are concepts that are fundamental to thought, or it may be because they refer to irreducible natural kinds. I will not speculate on the matter here. Instead, I offer one final note. Since both world-bound concepts and world-invariant concepts have a world-to-mind direction of determination, they refer to properties about which the non-indexical beliefs of every member of the linguistic community could be false or incomplete; but since Twin Earth thought experiments can be generated only for the former, it is clear that content externalism rests fundamentally not on the possibility of Twin Earth scenarios, but, as mentioned earlier, on the possibility of ignorance.

## 5. Temporal range

Just as concepts can differ both with respect to their direction of determination and with respect to their modal range, concepts can also differ with respect to their temporal range. Concepts fall into one of three kinds in this respect; concepts can be 'time-invariant', 'time-bound', or 'externally-anchored'. Any concept with a world-to-mind direction of determination is time-invariant in the specific, relatively weak sense that it transcends certain kinds of changes in our linguistic practices across time. This feature is a direct consequence of the fact that concepts with a world-to mind direction of determination are not determined ultimately by our conceptions, individual or communal, in contrast to time-bound and externally-anchored concepts, which are. Time-invariant concepts, being concepts with a world-to-mind direction of determination, were the focus of the previous section, and I will not discuss them further here. In this section, I focus instead on the distinction between time-bound concepts and externally-anchored concepts, where this

---

[19]The discussion of normative concepts in Eklund (2017) can itself be understood as addressing the question of the modal range of normative concepts.

distinction marks a division amongst concepts with a mind-to-world direction of determination.

Before explaining the distinction between time-bound and externally-anchored concepts in more detail, there are two important points to note. First, whereas both a concept's direction of determination and a concept's modal range are essential properties of it, a concept's temporal range is a contingent feature of it. To clarify, I take all of the key concepts I have introduced in this paper to be world-invariant; but to say that a concept is world-invariant is not to say that the having of the property referred to by that concept is essential to the identity of the concept in question. Second, concepts with a mind-to-world direction of determination lie on a continuum ranging from the time-bound at one end to the (increasingly) externally-anchored at the other. The two points are related. The extent to which a practice-bound concept is externally-anchored depends on the specifics of the institutionalised practices that both surround it and help to individuate it, and these practices can change over time. This means that exactly where a concept falls on the continuum is subject to change over time as an 'external anchor' can be introduced, tightened, or let slip. I will say more about external anchors shortly, but for now I provide some intuitive examples of concepts that fall into each category.

Concepts such as *game* and *banter* provide plausible examples of time-bound concepts. They are time-bound to the extent that they are determined by the conceptions (and hence linguistic practices) of the community at the time. Which concepts are expressed by such terms will therefore change as the communal conceptions and linguistic practices shift. The terms that express time-bound concepts are thus maximally sensitive to the vagaries of communal conceptions and linguistic practices; they are, more than any other kind of term, subject to semantic and conceptual drift. Such changes can happen at alarming rates, with terms sometimes displaying category-shift almost immediately, from noun to adjective to adverb to verb, and so on.

Concepts such as *criminal* and *Catholic bishop,* in contrast, provide plausible examples of externally-anchored concepts. The legal system serves as an external anchor for the former, since it is only relative to a legal system that one can be a criminal; and the institution of the Catholic Church serves as an external anchor for the latter, since it is only by being ordained by the Catholic Church that one can be a Catholic bishop. Although these concepts, being practice-bound, are determined ultimately by the conceptions of the community, the conceptions of the

community, as well as the related patterns of deference, are themselves in these cases anchored to and shaped by the relevant institutionalised practices. The institutionalised practices therefore provide a kind of stability to externally-anchored concepts that time-bound concepts do not have. In cases where institutionalised practices help to promote a just society, the stability provided by the external anchor will constitute a good; but in cases where institutionalised practices contribute to social injustice, the stability provided by the external anchor will constitute a harm. Since social institutions are typically complex entities, a single social institution will typically serve to anchor a wide range of concepts and will promote justice in some respects while hindering it in others.

A significant implication of the complex relationship between externally-anchored concepts and the institutionalised practices that both serve to anchor them and feature in relevant meaning-determining patterns of deference is that the conceptual engineering of such concepts necessitates the social engineering of the relevant institutionalised practices themselves. This is what marks out externally-anchored concepts for special ameliorative consideration.[20] This feature is nicely illustrated by the recent conceptual engineering of the concept *marriage,* which went hand-in-hand with the social engineering of the related institutionalised practices of marriage in the UK and elsewhere. Arguably, the concept would not have changed had there been no related change to the institutionalised practices themselves, where these changes were initiated by changes in the surrounding beliefs about what marriage ought to be. Similarly, no matter what proportion of the community believes that it should be possible for women to be ordained in the Catholic Church, the concept *Catholic bishop* will not change without changes to the related institutionalised practices of the Catholic Church itself.

External anchors differ in scope, strength and formality. Such differences have implications for the relative stability they provide to the concepts that are anchored to them, as well as for the potential goods or harms that result. Sometimes an external anchor can serve merely to anchor the concept a word expresses within a specific, relatively narrowly circumscribed context. Thus, when McVities successfully argued in a court of law that Jaffa cakes should be classified as cakes and not biscuits, the ruling, which served as an external anchoring for the concepts *cake* and *biscuit* had a relatively narrow scope, although it nonetheless had

---

[20]There is more to be said about the specific issues that externally-anchored concepts raise in connection with conceptual engineering, but I will have to leave this for another occasion.

financial implications for a significant number of people, given the different way in which cakes and biscuits can be taxed in the UK. Which mind-to-world concept is expressed by a term may therefore differ from context to context. Other external anchors, such as the various institutions and traditions of marriage across the world, as well as the various legal systems in different countries, have far-reaching implications for everyone. Finally, external anchors need not be formally recognised, legally-backed institutions; they can be informal, and they can be illegal.

As with world-to-mind concepts, the determination conditions for externally-anchored concepts are not to be confused with their acquisition conditions. An individual may acquire an externally-anchored concept by interacting with instances of the property to which the concept refers. But this does not mean that the concept has a world-to-mind direction of determination. To reiterate a point made earlier, concept acquisition is essentially an individual matter; concept determination, in contrast, is not essentially an individual matter. Externally-anchored concepts are concepts with a mind-to-world direction of determination no matter how acquired.

We can define what it is for a concept to be time-bound and what it is for a concept to be externally-anchored as follows:

(TBC) A time-bound concept is one which has a mind-to-world direction of determination and is determined ultimately by the communal conceptions at the time.

(EAC) An externally-anchored concept is one which has a mind-to-world direction of determination and is determined in part by its relations to a set of institutionalised practices that serve as an external anchor.

The principles are, once again, mutually exclusive and jointly exhaustive but silent on the actual classification of examples, and here too there is room for disagreement. I have argued that concepts can vary with respect to their direction of determination, their modal range and their temporal range. Now is the time to consider the implications of these variations for meta-semantic accounts of conceptual engineering.[21]

---

[21]I have said nothing in this paper about whether empty predicative terms express concepts, and if so, what kinds of concepts they express. Although I do not have the space to elaborate further here, there is reason to think that empty concepts can also vary according to their direction of determination, their modal range and their temporal range, and that this will also have implications for projects in conceptual engineering and for a meta-semantic framework able to accommodate them.

## 6. Alternative meta-semantic frameworks

Examples of concepts that have been cited as the focus of projects in conceptual engineering form a diverse group. They include: *belief, knowledge, truth, gender, race, good, right, marriage, person, immigrant, gene, fish*. It should be clear from what I've said above that these concepts are not all of the same kind. As a result, what is involved in conceptually engineering one such concept may be different from what is involved in conceptually engineering another. This means that, at a higher level of abstraction, a meta-semantic framework capable of accommodating the diverse array of projects in conceptual engineering will itself have to be sensitive to the different kinds of concepts that might be involved. Conceptually engineering an externally-anchored concept requires the social engineering of the institutionalised practices to which it is anchored; conceptually engineering a concept with a world-to-mind direction of determination, in contrast, does not—indeed, cannot—involve engineering aspects of the mind-independent world, such as the property to which a time-invariant concept refers. A central aim of this paper has been to draw attention to these differences. But it should be clear that the meta-semantic framework I've offered, distinguishing as it does between meanings, concepts and conceptions, is in a good position to accommodate them.[22] I now consider some alternatives.

First, consider the account of conceptual engineering offered in Cappelen (2018). According to Cappelen's 'austere' meta-semantic framework, conceptual engineering involves a change to the intensions and extensions of our terms but has nothing to do with concepts. In changing the extensions of the terms, we change the nature of the things our terms are about. For example, when we conceptually engineer terms such as 'family' and 'marriage', we change the nature of families and marriage respectively. This is why the account is dubbed a 'worldly' account of conceptual engineering. Cappelen's claims with respect to the conceptual engineering of terms such as 'family' and 'marriage' are highly plausible. Set within the taxonomy of concepts I have offered, this is because concepts such as *family* and *marriage* are externally-anchored, where, as noted above, the conceptual engineering of an externally-anchored concept necessitates the social engineering of the anchoring

---

[22]Distinguishing more clearly between the way in which world-bound concepts are externally-individuated and the way in which temporally-anchored concepts are externally-individuated would provide a solution to some of the problems encountered by Haslanger's appeal to 'objective type externalism' in the discussion of ameliorative projects concerning race and gender (cf. Haslanger 2006). I do not have the space to develop this line of thought fully here.

institutionalised practices. The problem for Cappelen's austere meta-semantic framework, however, is that every case of conceptual engineering has to be treated in the same way. This means that every instance of conceptual engineering has to be construed as a case of worldly conceptual engineering. The generalisation of the worldly account is a direct result of the austerity of the meta-semantic framework. But the generalisation is implausible. It is implausible to think, for example, that conceptual engineering could change the nature of truth or the nature of fish. Rather, conceptual engineering projects related to the terms 'truth' and 'fish' plausibly involve discovering something about the nature of truth and fish respectively. The implausibility can be traced, according to the taxonomy of concepts I've provided, to the fact that terms such as 'truth' and 'fish' express concepts with a world-to-mind direction of determination, where the meaning of the term, but not the concept expressed, is determined ultimately by our conceptions.

It might be thought that the implausible generalisation can be blocked by providing a metaphysically deflationary reading of what is involved in the worldly account of conceptual engineering. According to a metaphysically deflationary reading, conceptual engineering does not involve changing the world in any metaphysically strong sense, but instead involves no more than changing the intensions and extensions of our terms, and noting that this can nonetheless always be described at the object level. The problem with this reading is that even if we grant that it makes a worldly account of conceptual engineering for terms such as 'truth' and 'fish' plausible (although I am sceptical even of this), the metaphysically deflationary reading thereby fails to do justice to the metaphysically robust sense in which, when we conceptually engineer the terms 'family' and 'marriage', we change the nature of families and marriage. What we need is an account of the difference between terms such as 'family' and 'marriage' on the one hand and terms such as 'truth' and 'fish' on the other that explains why the metaphysically robust account of worldly conceptual engineering is appropriate for the former but the metaphysically deflationary reading of worldly conceptual engineering is appropriate for the latter. The austerity of the meta-semantic framework prevents us from drawing such a distinction.

The problem of implausible generalisation will affect any account of conceptual engineering which is both purely semantic and maintains that the intensions and extensions of the relevant terms always change. It will also, for similar reasons, affect any account of conceptual engineering according to which the phenomenon to which a given term or

concept refers is always engineered in the process. Sometimes conceptual engineering involves the engineering of the phenomenon to which the relevant term or concept refers; but sometimes it does not. The difference, I have argued, depends in part on the direction of determination of the concept in question. Finally, the problem of implausible generalisation will also affect deflationary accounts of conceptual engineering that construe conceptual engineering merely as conceptual analysis, since although conceptual engineering can sometimes be understood in this way (most notably for world-invariant concepts such as *truth, causation*, and *justice*), it cannot always be understood in this way (with externally-anchored concepts such as *marriage* and *criminal* providing obvious exceptions).[23]

Before closing, I will briefly comment on Pinder's claim that 'conceptual engineering … operates independently of metasemantics' (Pinder 2020, 18). The alleged separation of conceptual engineering from metasemantics follows from Pinder's claim that conceptual engineering operates at the level of speaker-meaning whereas meta-semantics operates at the level of semantic meaning. According to Pinder, it is 'sufficient for conceptual engineering that a term be explicitly used *as if* it has a newly constructed, non-standard definition' (Pinder 2020, 10, original emphasis) in a local context, where this is to be understood as providing the speaker-meaning of the term in that context, thereby leaving the semantic meaning of the term (and its intension and extension) unchanged. But this does not capture examples of conceptual engineering as generally understood. Using the term 'marriage' in a clearly-defined, non-standard way in a local context so as not to exclude same-sex couples did not count as conceptually engineering the term 'marriage'; a change to the institutionalised practices was also required (bringing along with it a change in semantic meaning and a change in the concept *marriage*). Similarly, using the term 'eligible voter' in a clearly-defined, non-standard way in a local context so as to include 16- and 17-year-olds does not count as conceptually engineering the term 'eligible voter'; more would be needed, such as communal uptake and, again, changes to institutionalised practices. Of course, Pinder is free to use the term 'conceptual engineering' explicitly as if it has a narrower definition that it is typically taken to have, but then we

---

[23]Cf. Deutsch (2020). See also Sawyer (2020b). The problem of implausible generalisation will also affect temporal externalism if construed as a general meta-semantic framework, since the claim that the meaning of a term is determined by future practice, even if plausible for certain cases (although I am sceptical even of this) is implausible for others. It is not clear, however, that temporal externalism is intended to provide a general meta-semantic framework. Cf. Ball (2020), following Jackman (1999, 2005).

would need to distinguish conceptual engineering in Pinder's sense from the broader issues that are of concern to conceptual engineers, where an understanding of these broader issues requires an appropriate meta-semantic framework, and in particular one that distinguishes different kinds of concepts.

The shift to a focus on speaker-meaning is motivated in part by an attempt to solve what is known as The Implementation Challenge. The thought runs roughly as follows: conceptual engineering involves meaning-change; but it's hard to see what we could do in order to effect a change in semantic meaning, especially if semantic meaning is externally-individuated; speaker-meaning, on the other hand, is, at least to some extent, under the control of the speaker; so, if we focus on speaker-meaning, we can solve the Implementation Challenge. But the examples of *marriage* and *eligible voter* given above illustrate that it is, contrary to what Pinder claims, *not* sufficient for conceptual engineering that a term be explicitly used *as if* it has a newly constructed, non-standard definition in a local context. Externally-anchored concepts cannot be conceptually engineered in this way precisely because of their relation to the institutionalised practices that serve to anchor them. But the motivation for shifting attention from semantic meaning to speaker-meaning in any case rests on a false presupposition. It rests on the false presupposition that because conceptual engineering *involves* meaning-change, it must *start* with meaning-change. On the contrary, in certain cases—notably those involving externally-anchored concepts—meaning-change plausibly occurs at the *end* of the process, not at the beginning. Think again of the marriage example. The initial aim in certain countries was to bring about a change in the institution of marriage so as to eradicate a certain kind of social injustice. Implementing the change involved campaigning to raise awareness, protesting and lobbying. None of these activities required any given speaker to use the term 'marriage' in a non-standard way, and a change in semantic meaning came about as a result of the institutional change, not vice versa. This means that The Implementation Challenge itself needs to be reconceived. Some projects in conceptual engineering will start with a change in speaker-meaning, but not all will. What is needed here too, then, is a sensitivity to the fact that concepts, and hence projects in conceptual engineering related to them, fall into distinct kinds.[24]

---

[24]Pinder also claims that 'there might be no substantive role in conceptual engineering for anything we should call "concepts"' (Pinder 2020, 3). However, since his focus is on conceptual engineering understood in a narrower sense than is typically understood, his discussion does not obviously connect with

Careful consideration of the wide variety of accounts of conceptual engineering that have been proposed falls outside the scope of the present paper. Notably, I have not considered the variety of functionalist accounts of conceptual engineering, in part because they cut across a range of different meta-semantic accounts, some being purely semantic, others embracing concepts, some providing a detailed account of a single case, others providing more general claims.[25] But the discussion nonetheless suggests that a proper understanding of conceptual engineering will need to be sensitive to the fact that there are different kinds of concepts, and that such sensitivity should be used as one of the measures against which any given meta-semantic framework is judged.

## 7. Conclusion

In this paper, I have argued that an adequate meta-semantic framework capable of accommodating the range of projects currently identified as projects in conceptual engineering must be sensitive to the fact that concepts (and hence projects relating to them) fall into distinct kinds. Concepts can vary, I have argued, with respect to their direction of determination, their modal range, and their temporal range. Acknowledging such variations yields a preliminary taxonomy of concepts and generates a meta-semantic framework that allows us both to accommodate the full range of cases and to identify a proper subset of concepts for special ameliorative consideration. Ignoring such variations, in contrast, leads to a restricted meta-semantic framework that accommodates only a subset of the particular projects while generating implausible accounts of others.

---

the issues I have raised in this paper. He does not engage with the meta-semantic framework I have offered, instead citing Cappelen's comment that people do not typically explain what concepts are. Interestingly, neither of the two assumptions Pinder claims are implicitly made about concepts are ones that I accept (cf. Pinder 2020, 3).

[25]See for example Brigandt (2010), Nado (2019), Prinzing (2018) and Thomasson (2020).

Conceptual Engineering Online Workshop hosted by ConceptLab (listed alphabetically by first name): Alex Horne, Angela O'Sullivan, Bernt Ivar Barkved, Cyrill Mamin, Davide Andrea Zappulli, Hsiang-Yun Chen, Jennifer Nado, Joey Pollock, Mark Pinder, Matheus Valente, Niklas Becker, Paul-Mikhail Catapang Podosky, Pedro Espirito Santo Abreu, Sigurd Jorem, Steffen Koch, Steven Diggin, and Tristram McPherson.

## Disclosure statement

## Funding

## References

Ball, D. 2020. "Revisionary Analysis Without Meaning Change (Or, Could Women be Analytically Oppressed?)." In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen, and D. Plunkett, 35–58. Oxford: Oxford University Press.

Brigandt, I. 2010. "The Epistemic Goal of a Concept: Accounting for the Rationality of Semantic Change and Variation." *Synthese* 177 (1): 1211–1241.

Burge, T. 1979a. "Individualism and the Mental." *Midwest Studies in Philosophy* 4: 73–121.

Burge, T. 1979b. "Sinning Against Frege." *Philosophical Review* 88: 398–432. Reprinted in Burge (2005).

Burge, T. 1986. "Intellectual Norms and Foundations of Mind." *Journal of Philosophy* 83 (12): 697–720.

Burge, T. 1990. "Frege on Sense and Linguistic Meaning." In *The Analytic Tradition*, edited by D. Bell, and N. Cooper, 30–60. Oxford: Oxford University Press. Reprinted with a postscript in Burge (2005).

Burge, T. 2005. *Truth, Thought, Reason: Essays on Frege*. Oxford: Oxford University Press.

Burge, T. 2014. *Origins of Objectivity*. Oxford: Oxford University Press.

Cappelen, H. 2018. *Fixing Language: An Essay on Conceptual Engineering*. Oxford: Oxford University Press.

Deutsch, M. 2020. "Conceptual Analysis Without Concepts." *Synthese*, doi:10.1007/s11229-020-02775-0.

Eklund, M. 2017. *Choosing Normative Concepts*. Oxford: Oxford University Press.

Haslanger, S. 2006. "Gender and Race: (What) Are They? (What) Do We Want Them to Be?" *Noûs* 34 (1): 31–55.

Horgan, T., and M. Timmons. 1991. "New Wave Moral Realism Meets Moral Twin Earth." *Journal of Philosophical Research* 16: 447–465.

Jackman, H. 1999. "We Live Forwards but Understand Backwards: Linguistic Practices and Future Behaviour." *Pacific Philosophical Quarterly* 80: 157–177.

Jackman, H. 2005. "Temporal Externalism, Deference, and our Ordinary Linguistic Practice." *Pacific Philosophical Quarterly* 86: 365–380.

Kripke, S. 1972. *Naming and Necessity*. Cambridge, Mass.: Harvard University Press.

Ludlow, P. 2014. *Living Words: Meaning Underdetermination and the Dynamic Lexicon*. Oxford: Oxford University Press.

Nado, J. 2019. "Conceptual Engineering, Truth and Efficacy." *Synthese*, doi:10.1007/s11229-019-02096-x.

Pinder, M. 2020. "Conceptual Engineering, Metasemantic Externalism and Speaker-Meaning." *Mind; A Quarterly Review of Psychology and Philosophy*, doi:10.1093/mind/fzz069.

Plunkett, D., and T. Sundell. 2020. "Metalinguistic Negotiation and Speaker Error." *Inquiry* 62 (1–2): 142–167.

Prinzing, M. 2018. "The Revisionist's Rubric: Conceptual Engineering and the Discontinuity Problem." *Inquiry* 61 (8): 854–880.

Putnam, H. 1970. "Meaning and Reference." *Journal of Philosophy* 70: 699–711.

Sawyer, S. 2007. "There is no Viable Notion of Narrow Content." In *Contemporary Debates in Philosophy of Mind*, edited by B. P. McLaughlin, and J. D. Cohen, 20–23. Oxford: Blackwell.

Sawyer, S. 2014. "Minds and Morals." *Philosophical Issues* 24 (1): 393–408.

Sawyer, S. 2018. "The Importance of Concepts." *Proceedings of the Aristotelian Society* 118 (2): 127–147.

Sawyer, S. 2020a. "Talk and Thought." In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen, and D. Plunkett, 379–395. Oxford: Oxford University Press.

Sawyer, S. 2020b. "Truth and Objectivity in Conceptual Engineering." *Inquiry* 63 (9-10): 1001–1022.

Sawyer, S. 2020c. "The Nature of Content: A Critique of Yli-Vakkuri and Hawthorne." *Inquiry.* doi:10.1080/0020174X.2020.1805702.

Sawyer, S. 2020d. "The Role of Concepts in Fixing Language." *Canadian Journal of Philosophy* 50 (5): 555–565.

Shields, M. 2020. "Conceptual Change in Perspective." *Inquiry* 63 (9-10): 930–958.

Thomasson, A. 2003. "Realism and Human Kinds." *Philosophy and Phenomenological Research* 67: 580–609.

Thomasson, A. 2020. "A Pragmatic Method for Normative Conceptual Work." In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen, and D. Plunkett, 435–458. Oxford: Oxford University Press.

Yli-Vakkuri, J., and J. Hawthorne. 2018. *Narrow Content*. Oxford: Oxford University Press.