

Talk and Thought

Sarah Sawyer

1. Introduction

Thought and language are uniformly acknowledged to be distinct phenomena requiring distinct philosophical treatment. Thought is one thing; language another. Nonetheless, they are intimately related. According to Frege, ‘The thought, in itself immaterial, clothes itself in the material garment of a sentence and thereby becomes comprehensible to us. We say a sentence expresses a thought’ (Frege 1918: 292). Abstracting away from Frege’s commitment to the immateriality of thoughts, the assumption that sentences express thoughts is widely accepted. It has, however, led to a near-ubiquitous identification between the linguistic meanings of non-indexical sentences on the one hand and the contents of the thoughts expressed by those sentences on the other. The identification is embedded in the widespread practice of appealing to a single proposition to specify both the semantic content of a non-indexical sentence and the content of the thought thereby expressed. Thus the semantic content of the sentence ‘Marriage is a legal union of two people’ is taken to be the proposition that marriage is a legal union of two people, which proposition is also taken to be the content of the thought the sentence expresses. At the level of words rather than sentences, the terms ‘linguistic meaning’ and ‘concept’ are typically treated as synonyms, as evidenced by the widespread practice amongst philosophers of switching freely between the two, often within a single sentence. What is said and what is thought are thus typically treated as identical.¹

But the identification of what is said and what is thought is, I will argue, a mistake. It involves a conflation of two distinct phenomena. In this chapter, I outline an externalist account of linguistic meaning and an externalist account of thought content that clearly distinguishes the two.² I advocate, instead, a dual-aspect theory of representation according to which sentences *have* semantic contents and *express* thoughts. Crucially, the two are not only theoretically distinct but diverge in actual

¹ The practice is so widespread that providing references would be an impossible task.

² The view and its implications for meaning-shift are presented in Sawyer (2018). In the current chapter I discuss a wider range of issues, focusing in particular on the nature and value of conceptual engineering. The distinction between linguistic meaning and thought content can be found in Burge (1986).

cases. Indeed, for creatures like us, divergence is the norm.³ Understood primarily as a thesis about terms rather than sentences, the externalist, dual-aspect theory of representation maintains that non-indexical linguistic terms such as ‘whale’, ‘number’, ‘explanation’, ‘fairness’, ‘marriage’, and ‘gender’ *have* linguistic meanings and *express* concepts. Linguistic meaning is here to be understood as linguistically encoded content rather than as reference or denotation; and concepts are to be understood as representational constituents of thoughts individuated at the level of sense rather than reference.⁴ Here too, the crucial point is that the linguistic meaning of a non-indexical term and the concept expressed by that term are not only theoretically distinct but diverge in actual cases. This is because, on the account I offer, linguistic meanings and concepts track different phenomena and play different explanatory roles.

The distinction, understood along the lines I propose, brings theoretical gains in a cluster of related areas. It provides an adequate account of meaning change across time, it accounts for the possibility of substantive agreement and disagreement across different theoretical frameworks, it accommodates at face value the phenomenon of contested meanings, and it explains both the nature and the value of conceptual engineering, placing the phenomenon of conceptual engineering in a framework of theoretical change more broadly understood in a way that addresses recent prominent concerns.

A caveat before we begin. The current chapter is focused on non-indexical terms only. I set aside discussion of indexical terms, such as ‘I’, ‘here’, and ‘now’ as well as of proper names, such as ‘London’ and ‘Sarah’, which I have argued elsewhere should be understood as containing an indexical element in their singular use.⁵ Indexical terms require separate treatment.⁶ I focus instead on the claim that the linguistic meaning of a non-indexical term should be distinguished from the concept expressed by that term. This is where the theoretical gains lie.

2. The Background: Internalism and Externalism

The (problematic) identification of the linguistic meaning of a non-indexical term and the concept expressed by that term sits most naturally within a thorough-going internalist theory of both thought and language. Internalism about thought and language are defined by their commitment to a local supervenience thesis according to which what a subject thinks and what she means by her words are each determined by her intrinsic (typically, physical) states. Internalist theories of thought and

³ As will become clear, narrowing the divergence between the two is the upshot of successful inquiry. See section 6 below.

⁴ See Frege (1892).

⁵ For my view on proper names, see Sawyer (2010), which has its roots in Burge (1973).

⁶ For present purposes, I take as an indexical term any term which requires a contextual application in order to determine a referent. That is, roughly, any term that admits of a content/character distinction in the sense of Kaplan (1989). Indexical terms do not obviously express concepts—or, at least, are not guaranteed to express a single concept across different occasions of use; and the semantic content of a sentence containing an indexical term is not obviously identical to the content of the thought expressed. Which terms fall into this category is a matter of debate.

language are designed to capture the way the world seems from the individual's perspective narrowly construed.⁷ Concepts and linguistic meanings, then, understood as the internalist understands them, are determined by the same set of intrinsic facts and are invoked to capture a single phenomenon, namely the subject's individual, perceptual, discriminatory capacities and inferential dispositions which together inform her deployment of a concept and her use of a term. This means that within a thorough-going internalist framework, not only is there no reason to distinguish the linguistic meaning of a term from the concept expressed by that term, there is also no means of doing so. It is precisely this feature of a thorough-going internalist theory that, I maintain, precludes it from providing an adequate account of phenomena such as meaning change, substantive agreement and disagreement across different theoretical frameworks, and the nature and value of conceptual engineering. I return to this issue in section 6 below.

A distinction can be drawn between the linguistic meaning of a term and the concept expressed by that term only within a framework that distinguishes between the facts that determine the former and the facts that determine the latter. A thorough-going internalist framework does not have the resources to do this, but a thorough-going externalist framework does.⁸ In fact, Putnam's original introduction of externalism to the mainstream literature provides the means to distinguish the two, but his exclusive focus on language obscures this important insight. Putnam, concerned with language rather than thought, identifies two elements that he claims are missing from traditional semantic theories.⁹ The first is the contribution of society, and the second is the contribution of the real world. The contribution of society is illustrated by Putnam's example involving his incomplete understanding of the terms 'beech' and 'elm', and brings in the core notion of linguistic deference, which he articulates as deference to 'experts'. The contribution of the real world is illustrated by Putnam's notorious Twin Earth thought experiment involving the term 'water', and brings in the core notion of causal relations to natural kinds with hidden essences. My suggestion, elaborated throughout the chapter, is that the contribution of society to representation is best understood as a contribution to language and the contribution of the real world to representation is best understood as a contribution to thought.¹⁰

It is reasonable to think not only that the contributions Putnam identifies can be understood as contributions to different phenomena, but, in addition, that they

⁷ See for example Fodor (1980, 1987), Stalnaker (1990), Segal (2000), and Chalmers (2003). For an internalist account of linguistic meaning most relevant to the concerns of the present chapter, see Sundell (2011, 2012) and Plunkett and Sundell (2013).

⁸ 'Mixed' theories which combine an internalist theory of thought and an externalist theory of language (or vice versa) also have the resources to do this; indeed, they necessarily entail a distinction between the linguistic meaning of a term and the concept expressed by that term, although this implication has not generally been explored. I do not discuss such theories in the chapter because the way in which they draw the distinction does not imply the theoretical advantages with which I'm concerned. For examples of mixed theories of the first kind see Putnam (1973, 1975) and Crane (1991). I know of no-one who holds a mixed theory of the second kind and can see little motivation for such a view.

⁹ See Putnam (1973, 1975).

¹⁰ As will become clear, I take the contribution of the real world to extend far beyond what is given by causal relations to natural kinds with hidden essences.

cannot be seen as contributions to a single phenomenon, whether that be to language, as Putnam maintains, or to thought. This is because there is a tension between the two contributions, a tension which is masked by an ambiguity in the notion of an 'expert'. It is standard in the externalist literature to assume that an expert is someone who is knowledgeable about the relevant subject matter. On this understanding, the contribution of society and the contribution of the real world would necessarily be aligned. On this view, the linguistic meaning of an individual's term is determined either by causal relations to the real world, or by deference to experts, the linguistic meaning of whose term is determined by causal relations to the real world. Either way, linguistic meaning is determined at root by relations to the real world; the only question is whether the determination relation is direct or indirect.¹¹

But the claim that linguistic deference involves deference to those who are knowledgeable about the relevant subject matter is an idealization, and one which cannot play the central role required of it in an account of *actual* linguistic deference. Linguistic deference must be deference to other members of one's linguistic community, and these may not be experts in the sense of being knowledgeable about the relevant subject matter. The history of science is the history of experts who were wrong, sometimes significantly so, about the subject matter that fell into their area of expertise. It would be wrong to conclude that they were not experts after all; on the contrary, to be an expert is to be worthy of deference despite the possibility of error. This fits with our actual linguistic and sociological practice surrounding the term 'expert'.

We must reject the standard externalist assumption, then, that an expert, for the purposes of linguistic deference, is necessarily someone who is knowledgeable about the relevant subject matter. Rather, an expert is, roughly, someone to whom others defer not because they *are* knowledgeable but because they are *perceived* to be so. Being perceived to be knowledgeable does not, of course, preclude being knowledgeable; but being knowledgeable is neither necessary nor sufficient for being perceived to be so. Note, however, that on this more realistic, less idealized understanding of an expert, linguistic deference provides no guarantee that the linguistic meanings of our terms hook up with the real world in the way that is supposed to be secured by direct causal relations. This is because the way the experts take the world to be may be different from the way the world is.¹² The contribution of society and the contribution of the real world will therefore sometimes take us in different directions. To resolve the tension, the two contributions that Putnam identifies not only can, but should be seen as contributions to different phenomena.

The contribution of society to representation, I suggest, is to be understood as a contribution to language, and the contribution of the real world to representation is,

¹¹ It is a mistake to think that the contribution of the real world is relevant to natural kind terms whereas the contribution of society is relevant to non-natural kind terms. Drawing the distinction along these lines is clearly not true to Putnam's original discussion which uses examples of natural kind terms such as 'beech' and 'elm' to illustrate the phenomenon of linguistic deference. It also fails to do justice to the hidden depths of the natural world, the reality of the non-natural world and the fundamental role of linguistic deference in language-learning.

¹² Errors may be both theoretical, in the form of false beliefs, and practical, in the form of misapplications of words and incorrect deployment of concepts.

I suggest, to be understood as a contribution to thought. These distinct ‘external’ facts ground a theoretically significant distinction between the linguistic meaning of a term on the one hand and the concept expressed by that term on the other. I take the distinction to be general, applying to terms across a broad range of disciplines and areas of inquiry. The distinction is not restricted to natural kind terms, or to empirical terms more broadly construed, but extends beyond the empirical realm to mathematical, logical and philosophical terms as well as to normative terms, including social and ethical terms. In what follows, I provide an externalist account of linguistic meaning and an externalist account of concepts, and I demonstrate some of the theoretical gains of distinguishing the two.

3. Linguistic Meaning

I take linguistic meaning to supervene on use. I reject the internalist claim, however, that the linguistic meaning of a term as used by a given individual depends solely on her use of the term. There is often widespread variance in the use of a term by different individuals across a community, but this variance is consistent with a term’s having a single meaning for all. Linguistic meaning in this communal sense is what dictionaries aim to record. They achieve this aim, in so far as they do, not by recording statistical averages across individual usage but by paying attention to patterns of usage and deference across the community. Patterns of deference reflect our general recognition of the fact that, for any given term, some individuals are more competent in its use than others. It is the use of the most competent together with patterns of deference amongst all (including the most competent) that determines linguistic meaning.

Linguistic deference is a much more general phenomenon than is sometimes recognized, applying to most, probably all, terms in the language. This includes, for example, very basic terms such as ‘red’, ‘table’, and ‘sharp’, and slang terms such as ‘wicked’, ‘newb’, and ‘banterous’. It is the generality of the phenomenon of linguistic deference that makes talk of expertise misleading. Not only does the term ‘expert’ carry connotations of knowledgeable status, as noted in the previous section, but it also carries connotations of a restricted application to scientific, theoretical, or technical terms only. This restriction is artificial, since the fundamentality of linguistic deference is built into the very nature of language-learning. Talk of competence is preferable, even if not perfect.

What I have said so far is couched at a relatively general level. At a more specific level, the linguistic meaning of a term at a time can be understood as the characterization of the relevant subject matter that members of the linguistic community would settle on at that time were they to reach reflective equilibrium in the context of a dialectic.¹³

The dialectic, for these purposes, is an honest, open debate, shorn of all subjective elements, in which participants aim for a characterization of the subject matter

¹³ For an account of linguistic meaning along these lines, see Burge (1986, 1989) and Sawyer (2007, 2018).

through reason and reflection on actual and hypothetical cases, deferring to the most competent as and when appropriate. The dialectic is not to be understood as involving maximal reflection on the subject matter, as this would inevitably extend the discussion beyond actual use to what was perceived to be ideal future use. Rather, the relevant notion is to be understood as full reflection within actual empirical and theoretical boundaries, where this is consistent with the participants agreeing that the subject matter has not yet been fully characterized. The primary focus of the dialectic is the characterization of a subject matter rather than the characterization of meaning *per se* because although the questions ‘What does ‘x’ mean?’ and ‘What is an x?’ are different questions, the former can be, and typically is, answered by answering the latter.¹⁴

Within the set empirical and theoretical boundaries, there may be terms for which no agreement would be reached. The account implies that for such terms there is no settled linguistic meaning at the time. I take the claim that at least some terms have no settled linguistic meaning at a time to be a natural implication of the fact that meaning depends on use, which is subject to flux and changes over time. I therefore see the implication that not every term will have a settled linguistic meaning as a virtue of the account. Language is organic, and this fact must be captured by an adequate theory of linguistic meaning.

Of particular significance in this context are terms with so-called ‘contested meanings’.¹⁵ The literature on conceptual engineering provides a set of staple examples, including ‘race’, ‘gender’, ‘marriage’, and ‘rape’, although terms with contested meanings occur across all realms of inquiry. The meanings of such terms are contested in the sense that different groups of people within the linguistic community—the conservative group and the progressive group, as we might call them—disagree in fundamental ways about how the term ought to be used, where this disagreement is reflected by the different actual uses of the relevant term by the different sub-sections of the community. In such cases, the community as a whole is aware of the disagreement in use, but the disagreement takes place against the common understanding that the matter would not be resolved by a stipulative disambiguation introduced to accommodate the different uses. The different parties to the dispute, each convinced that their own use is correct, do not see themselves as talking past each other, but as offering different views on a single subject matter. In the context of a dialectic, even if each sub-section of the community were to settle on an agreed characterization of the relevant subject matter, no overall agreement would be reached by the community as a whole. The fundamental disagreement over the correct characterization of the subject matter in such cases splits the use of the relevant term and renders its meaning contested. I return to these important cases in section 5 below.

The account of linguistic meaning offered accommodates the fact that the meaning of a term can change over time. Since meaning supervenes on use, a change in the linguistic meaning of a term depends on an underlying change in linguistic practice;

¹⁴ This point is made in Burge (1986), where it is related to the claim made in Quine (1951) that there is no separating truths of meaning from matters of fact.

¹⁵ See Gallie (1956) for an early discussion of ‘essentially contested meanings’.

and if there is a change in linguistic practice, there will be a change in the characterization that would be settled on in the context of a dialectic. For example, the meaning of the term 'meat' has clearly changed between Shakespearean times and now. The term 'meat' used to mean something like *food in general*, whereas now it means something like *animal flesh that is eaten for food*. But this is precisely the result that would emerge in the context of dialectical reflection on use at the two times.

Given that the linguistic meaning of a term determines its extension, the account of linguistic meaning also accommodates the fact that the extension of a term can change over time. I take the extension of a term to be the class of entities that satisfy the term's (descriptive) linguistic meaning. Thus it is because an apple satisfies the description 'food in general' that it falls into the extension of the term 'meat' in Shakespearean times; and it is because an apple does not satisfy the description 'animal flesh that is eaten as food' that it does not fall into the extension of the term 'meat' now. Finally, the account also provides linguistic norms by establishing how the relevant term ought to be used, namely in accordance with the agreed characterization of the relevant subject matter. I return to questions of normativity and truth in section 7 below.

4. Concepts

A quick internet search provides an extensive list of words that have changed their meaning over time, including 'meat', 'spinster', 'bachelor', 'clue', 'awesome', 'awful', 'wicked', 'girl', 'egregious', 'pretty', and 'hussy'. In all of these cases, and many more besides, the change in meaning has been accompanied by a change in topic, or subject matter.¹⁶ This is what makes such cases relatively unproblematic from a philosophical perspective.

The more interesting cases are those for which we want to say that the meaning of the relevant term has changed while the subject matter has not. Such cases are widely regarded as philosophically problematic. To illustrate the nature of the problem raised by this kind of case, I start with a puzzle articulated by Sainsbury, although a variation of the puzzle occurs across a wide range of philosophical literature.¹⁷ Let us assume, as Sainsbury does, that the term 'whale' was embedded in a linguistic practice in ancient times, when people thought whales were fish, that is different from the linguistic practice in which it is embedded now, when people think whales are mammals. This, he says, raises a dilemma. Either the sentence 'Whales are fish' means the same in ancient times as now, or it doesn't. To say that it does fails to accommodate the fact that meaning is determined by use; but to say that it doesn't fails to accommodate the fact that there is substantive disagreement across the two times.

Substantive disagreement is to be contrasted with merely verbal disagreement, such as the kind of disagreement that might occur over whether an apple is a form of meat between someone from Shakespeare's time and someone from the present time,

¹⁶ Unlike technical terms such as 'extension' and 'reference', I take the terms 'subject matter' and 'topic' to be non-technical and relatively intuitive.

¹⁷ For Sainsbury's articulation of the puzzle, see his (2014).

the former insisting that an apple is a form of meat, the latter insisting that it isn't.¹⁸ The disagreement would be merely verbal in the sense that the two parties would be talking past each other. One way to capture this is to note that there is no single content over which the parties disagree. We can stipulate that 'meat_s' is to mean what 'meat' means in Shakespeare's time, and 'meat_p' is to mean what 'meat' means in the present. It is plausible to assume that the disputants, once apprised of the difference in use, would accept the stipulations and agree that apples are a form of meat_s but not a form of meat_p. The dispute would not persist after the stipulation because it would be clear that the subject matter of the term 'meat' in Shakespeare's time is not the same as the subject matter of the term 'meat' now. Disambiguation works in this case because it separates what are clearly two distinct subject matters.

Substantive disagreement, in contrast, is disagreement over a single subject matter. The disambiguation strategy that works in cases of merely verbal disagreement does not, therefore, work in cases of substantive disagreement. The disagreement over the truth of the sentence 'Whales are fish' persists even if we stipulate that 'whale_a' is to mean what 'whale' means in ancient times and that 'whale_p' is to mean what 'whale' means in the present. This is because the stipulation does not disambiguate two distinct subject matters; rather, it distinguishes two theories about a single subject matter. The question remains which theory, if either, is correct. The disagreement is fundamentally a disagreement over the nature of whales.¹⁹

Sainsbury's puzzle is puzzling, then, because, intuitively, we want to be able to say both that the linguistic meaning of the sentence 'Whales are fish' is different at the two times and that the subject matter of the sentence is the same at the two times, and it is unclear exactly how we can do both. The reason it is unclear how we can do both is that sameness of subject matter appears to require a single propositional content over the truth of which the parties to the dispute disagree, and this appears to require sameness of linguistic meaning which directly contradicts the claim that the linguistic meaning of the term has changed.

The diagnosis of the puzzle lies in the recognition that if linguistic meaning is the only representational element in our theory, it is subject to inconsistent constraints. Linguistic meaning cannot both supervene on use *and* determine a stable subject matter. This is because in order to supervene on use, linguistic meaning must change in accordance with a change in linguistic practice that comes about as a result of a change in the community's beliefs; but in order to determine a stable subject matter, linguistic meaning must be insensitive to at least some changes in the community's beliefs and hence insensitive to at least some changes in linguistic practice. No single element can do both.

¹⁸ On the nature of verbal disputes, see for example Chalmers (2011). Although I agree with much of what Chalmers says, I think the account ultimately suffers from the conflation between linguistic meanings and concepts that I urge in the current chapter. I do not have the space to discuss Chalmers's views in detail here.

¹⁹ The boundary between merely verbal disagreements and substantive disagreements may be vague and will certainly sometimes appear so, given that the distinction between subject matters is not necessarily transparent. This is consistent, however, with there being clear-cut cases on either side of the boundary.

The solution to the puzzle, then, begins with the recognition that two representational elements are required; one to supervene on use, and the other to determine a stable subject matter. Of the two roles, linguistic meaning is ideally suited for the former. The account of linguistic meaning offered above starts from the assumption that meaning supervenes on use, and explains how a change in linguistic meaning tracks changes in a community's linguistic practice. The second role is that of determining a stable subject matter. This, I maintain, is the function of concepts. Concepts are constituent, representational elements of thoughts that connect thinkers representationally to a subject matter about which individual or communal beliefs may vary.

In order for concepts to play the requisite role of securing a stable subject matter, they must be understood as externally individuated, fundamentally non-descriptive components of thought. They must be externally individuated if they are to determine a subject matter that can be stable across individuals with different individual beliefs; and they must be fundamentally non-descriptive if they are to determine a subject matter that can be stable across communities with different communal beliefs. Concepts are not individuated by individual conceptions—they are not individuated by the way the individual thinker takes the world to be. Nor are they individuated by communal conceptions—they are not individuated by the way the community as a whole takes the world to be. Concepts are individuated, at the fundamental level, by relations to objective properties. This is the sense in which the contribution of the real world to representation is best understood as a contribution to thought. It is the fact that the subject matter itself enters into the individuation conditions of the relevant concept that explains how the concept expressed by a term can determine a stable subject matter.²⁰

Once two distinct representational elements are acknowledged, Sainsbury's puzzle is resolved. We wanted to be able to say both that the linguistic meaning of the sentence 'Whales are fish' is different at the two times, and that the subject matter of the sentence is the same at the two times. This we can now do. The linguistic meaning of the term 'whale' has changed over time, but it expresses the same concept at the two times, and hence concerns the same subject matter; the linguistic meaning of the sentence has changed over time, but it expresses the same thought at the two times, thereby securing a single propositional content over the truth of which the communities can be understood to have a substantive disagreement.

The assumption that substantive disagreement presupposes a single propositional content over the truth of which the parties to the dispute disagree is almost ubiquitous. It is questioned by Plunkett and Sundell, who argue that intuitions about substantive disagreement can be accommodated without appeal to a shared propositional content.²¹ It is interesting to note that both sides to this particular dispute assume that there is only one propositional content in question, which is both the

²⁰ I say 'at the fundamental level' to allow that there may be some descriptive concepts and some empty concepts. By 'fundamental', I mean representationally fundamental, and I take a liberal view on the issue and include amongst the fundamental concepts ordinary concepts such as those expressed by the terms 'whale', 'marriage', 'race', 'gender', 'belief', 'rape', 'moral goodness', and 'justice'.

²¹ See Sundell (2011, 2012) and Plunkett and Sundell (2013).

semantic content of the sentence and the thought expressed by it. The mainstream assumes that this propositional content must be identical in cases of substantive dispute; Plunkett and Sundell argue that it need not be. I offer a middle way. I agree with the mainstream that the best explanation of our intuitions about substantive disagreement is that there is a single propositional thought content over the truth of which the parties to the dispute disagree; I agree with Plunkett and Sundell that substantive disagreement does not require sameness of linguistic meaning.

5. Theoretical Frameworks and Contested Meanings

Distinguishing the linguistic meaning of a term from the concept expressed by that term in the way that I have suggested offers a solution to the question of how substantive disagreement can occur across theoretical divides. The linguistic meaning of a term at a time is determined by the received theory of the relevant subject matter at that time; but the concept expressed by a term is not determined by the theoretical framework within which it occurs; rather, the concept expressed by a term provides an anchor to a shared, objective world. This means that concepts may be shared across different communities who hold substantially different theories of the same subject matter. The account thus avoids a pernicious kind of conceptual relativism and overcomes the kind of incommensurability that Kuhn thought plagued theory change.²²

For the same reason, it provides a proper understanding of the phenomenon of contested meanings. As noted in section 3 above, we can say that the meaning of a term is contested when there is disagreement in use between sub-sections of the community and the disagreement would not be resolved by a stipulative disambiguation introduced to accommodate the different uses because the disagreement concerns the correct characterization of a single subject matter. In this sense, contested meanings essentially involve substantive rather than merely verbal disagreements. The term ‘marriage’ has a contested meaning in this sense. Use of the term ‘marriage’ by conservatives differs from use of the term ‘marriage’ by progressives, the former insisting that marriage is necessarily a union of a man and a woman, the latter disagreeing.²³ But the disagreement would not be resolved by a stipulative disambiguation, with the meaning of ‘marriage_c’ determined by the conservative use and the meaning of ‘marriage_p’ determined by the progressive use. It would not be resolved because the progressives’ very point is that the sex of the individuals concerned is irrelevant to the question of marriage—that there are not two kinds of marriage, but one. The disagreement, then, is a substantive disagreement about the nature of marriage. The stipulation does not disambiguate two distinct subject matters; rather, it distinguishes two theories about a single subject matter.

The important point to note here is that the characterization of contested meanings makes essential reference to the fact that they involve substantive disagreements about a single subject matter. This incurs an explanatory debt. We need an account of

²² See Kuhn (1962).

²³ The relevant issue in this context is not about whether the law precludes same-sex couples from being married but about whether there is something about the nature of marriage that precludes it.

how it can be that a term with a contested meaning is about a single subject matter when the parties to the disagreement use the term in such different ways. The theory of concepts offered in the previous section provides the requisite explanation. What makes the disagreement between conservatives and progressives a disagreement about marriage is that the term ‘marriage’, despite having a contested meaning in the community at large, expresses a single concept that anchors the use of the term by both groups to a single subject matter. The distinction between linguistic meanings and concepts I have suggested, then, provides an adequate explanation of the phenomenon of contested meaning by explaining how a disagreement in use between sub-sections of society can be a substantive disagreement about a single subject matter. Which terms have contested meanings will naturally change over time as a settlement is reached for some terms and disagreement arises over others. Contested meanings are grounded in theoretical differences within a community at a time. The changing landscape of contested meanings reflects the forces of theoretical and social change, and is an important phenomenon for this very reason.

6. Conceptual Engineering

The phenomenon of contested meanings has an important role to play in the context of conceptual engineering, which also faces a number of problems that can be resolved by distinguishing between linguistic meanings and concepts in the way that I have suggested. Let us take as an example of a recent project in conceptual engineering Haslanger’s proposed revisionary analysis of ‘woman’.²⁴ According to Haslanger:

S is a woman iff

- i. S is regularly and for the most part observed or imagined to have certain bodily features presumed to be evidence of a female’s biological role in reproduction;
- ii. that S has these features marks S within the dominant ideology of S’s society as someone who ought to occupy certain kinds of social position that are in fact subordinate (and so motivates and justifies S’s occupying such a position); and
- iii. the fact that S satisfies (i) and (ii) plays a role in S’s systematic subordination, that is, *along some dimension*, S’s social position is oppressive, and S’s satisfying (i) and (ii) plays a role in that dimension of subordination (2012: 234, original emphasis).

Revisionary analyses of this kind must satisfy two constraints: they must be both revisionary and conservative. They must be revisionary in the sense of proposing a new meaning for the relevant term, where this typically has the effect of shifting its extension; and they must be conservative in the sense of being an analysis of the

²⁴ Haslanger uses the term ‘ameliorative’ and there is a question about whether and in what way an ameliorative analysis is to be understood as revisionary. I take the distinction between linguistic meanings and concepts to help clarify this issue too, but do not have the space to discuss it here. I use the term ‘revisionary’ because of its general applicability.

original subject matter. If they do not satisfy the former constraint they are merely descriptive projects; and if they do not satisfy the latter constraint they have not provided an analysis of the relevant subject matter but have simply changed the topic. This creates a puzzle, since it is unclear how both constraints can be satisfied at once. The problem is that a change in extension brought about by a change in meaning appears to result in a change in topic; thus satisfaction of the first constraint appears to preclude satisfaction of the second.²⁵

The puzzle is driven by the (false) assumption that a change in extension corresponds to a change in subject matter, or topic. The assumption fits naturally within the kind of view that identifies linguistic meanings and concepts—the kind of view I have been arguing against. On such a view, there is a single representational element that determines both extension and subject matter, and hence it is possible (and natural) to identify (or conflate) the two. But the solution to the puzzle requires their separation.²⁶ One of the theoretical benefits of the dual-aspect theory of representation I am advocating is that it not only accommodates the separation of extension from subject matter, but it actually implies that the two are distinct and, moreover, that they can easily diverge. According to the account of linguistic meaning provided in section 3, the linguistic meaning of a term determines an extension descriptively. According to the account of concepts provided in section 4, the concept expressed by a term determines a subject matter non-descriptively. This means, effectively, that the extension of a term is determined by theory, whilst the concept expressed by a term is not. As a result, concepts can provide an anchor to a stable subject matter, or topic, about which there may be different proposed analyses.

Consider this in the context of Haslanger's revisionary analysis of 'woman'. The account of linguistic meaning I have offered implies that the extension of the term 'woman' will change if the traditional meaning of the term is replaced by the proposed revisionary analysis. Not all women in the traditional sense are women in Haslanger's sense. But the account of concepts I have offered does not imply that the subject matter will change as a result. This is because it is consistent with a change in the extension of the term 'woman', underwritten by a change in linguistic practice, that the concept expressed by the term nonetheless secures the same subject matter: women. Indeed, I take this to be the most plausible understanding of what is going on in this and many other cases of conceptual engineering and conceptual ethics. My view applies equally well, for example, to debates surrounding the nature of race, gender, marriage, consent, property, rape, and personhood, each of which has a history of revisionary analyses grounded in substantive disagreement.²⁷

²⁵ The concern is reminiscent of Strawson's objection to Carnap's views on conceptual explication on the grounds that conceptual explication implies a change in extension which amounts to a change in topic. See Carnap (1947) and Strawson (1963). My proposal is also, then, a response to Strawson on behalf of Carnap.

²⁶ This is recognized by Cappelen, who provides an alternative way to distinguish topic and extension. See Cappelen (2018). A discussion of the relative merits of the two approaches will have to await discussion on another occasion.

²⁷ For revisionary analyses of race, see for example Appiah (1992) and Haslanger (2012). See Clark and Chalmers (1998) for a revisionary analysis of belief.

In this sense, a revisionary analysis is not intended to describe our current linguistic practice but is intended as the analysis of the relevant subject matter that we *ought* to accept. This is because a proposed revisionary analysis is, in well-intentioned cases, an attempt, by those who take themselves to have more clearly conceptualized the subject matter, to initiate positive theoretical or social change.²⁸ In effect, the analysis contributes a revised characterization of the subject matter to the dialectic, thereby altering the status of the relevant term from one with a (relatively) stable meaning to one with a contested meaning, grounded in substantive disagreement about a subject matter. If the revisionary analysis is correct and accepted, the effect is to bring the extension of the linguistic meaning of a term in line with the extension of the concept it expresses (i.e., in line with the relevant subject matter); it moves linguistic practice closer to the truth. The analysis has to be both true and accepted to achieve this aim, since an unaccepted truth would not change linguistic practice, and an accepted falsehood would not bring linguistic practice closer to the truth. Conceptual engineering is, at its most interesting, the result of an attempt to uncover facts that we are partially aware of but have not yet fully grasped, whether mathematical, logical, philosophical, natural, social, or moral.²⁹ This does justice to the kind of normativity that underlies conceptual engineering projects by connecting it to questions about the way we ought to think and talk, as well as to questions about how we ought to act.³⁰

Burgess and Plunkett have raised a general concern for externalist accounts of conceptual engineering, responding to which will help to clarify my view. They write:

The textbook externalist thinks that our social and natural environments serve as heavy anchors, so to speak, for the interpretation of our individual thought and talk. The internalist, by contrast, grants us a greater degree of conceptual autonomy. One salient upshot of this disagreement is that effecting conceptual change looks comparatively easy from an internalist perspective. We can revise, eliminate, or replace our concepts without worrying what the experts are up to, or what happens to be coming out of our taps. (2013a: 1096)

The objection assumes that conceptual engineering involves revising, eliminating, or replacing our concepts. But conceptual engineering as I understand it does not involve revising our concepts. Indeed, there is an incoherence to the suggestion that it does, both for the externalist and the internalist. So long as we assume that concepts have their intensions and extensions essentially, it will not be possible for them to be revised. Nor does conceptual engineering involve eliminating or replacing our concepts. This would, I take it, amount to changing the topic of inquiry, which would not do justice to the interest and importance of the relevant projects in conceptual engineering. My account, of course, has two representational elements:

²⁸ I say 'in well-intentioned cases' because it is clearly possible for someone to propose, with evil intentions and for personal gain, a revised analysis which knowingly subverts the truth. The idea is explored through Orwell's use of the language 'Newspeak' in his (1949) novel *Nineteen Eighty-Four*. I restrict my focus to the kinds of examples I mention in the chapter, which, I take it, are all well-intentioned.

²⁹ For a mathematical example, see Frege's revisionary analysis of number in his (1884).

³⁰ The importance of the normative aspect of conceptual engineering is emphasized in Burgess and Plunkett (2013a,b).

concepts and linguistic meanings. And although conceptual engineering does not involve revising, eliminating, or replacing the former, it does involve revising, eliminating, and replacing the latter. In this sense, Cappelen is right when he says that conceptual engineering is about ‘fixing language’; it is, as he says, about better ways of talking about a topic.³¹ But there are two points to note here. First, talking about a topic requires a representational relation between us and the topic, and one which does not vary with the variation in our ways of talking. This is provided, I have argued, by concepts that are individuated in part by relations to objective properties in the world beyond us. Externally individuated concepts, then, provide the stable background against which conceptual engineering in the form of linguistic, and hence theoretical, change can take place. This means that a representational anchor to the world is an advantage of the account I have offered, and one which internalist theories cannot provide.³² Second, conceptual engineering as I understand it is a form of theorizing. This presents no specific problem for the externalist account I have suggested, which incorporates a notion of linguistic meaning that floats free from the heavy anchors of the social and natural environments accepted by the ‘textbook externalists’. A revisionary analysis, which can be proposed by any member of a linguistic community, will inevitably, in virtue of its revisionary status, go against the mainstream ‘expert’ opinion. The difficulty lies in persuading the mainstream of the merits of the revisionary analysis, but this is a difficulty that must be faced no matter which theory of representation is true.

7. Truth and Normativity

Having distinguished the linguistic meaning of a term from the concept expressed by that term, and hence the semantic content of a sentence from the thought expressed by that sentence, we need to separate the elements that are appropriate to language and the elements that are appropriate to thought.

One particularly important such property is truth.³³ I take truth, ultimately, to be a property of thought. The truth-value of a sentence, then, depends on the truth-value of the thought it expresses rather than on its (descriptive) semantic content. The two representational elements are governed by distinct norms: conceptual norms (norms of thought), and linguistic norms (norms of language). Conceptual norms concern truth. Linguistic norms, as stated in section 3 above, concern use in accordance with linguistic meaning at a time. For example, suppose that rape was standardly defined as possible only outside of marriage in 1800, but was standardly recognized (correctly) to be possible within marriage in 2000.³⁴ The sentence ‘Rape can occur in marriage’, then, expressed the same, true thought in 1800 as it did in 2000. But an utterance of the sentence in 1800 would have violated the linguistic norms of the time, whereas an utterance of the sentence in 2000 would have conformed to the

³¹ Cappelen (2018).

³² A general argument against internalism on these grounds is given in Sawyer (2007).

³³ For these purposes I do not take any specific stand on the nature of truth.

³⁴ For a history of substantive disagreement concerning rape, see Hasaday (2000).

linguistic norms of the time. This captures one aspect of the revisionary nature of conceptual engineering. A proposal that violates linguistic norms is bound to be regarded as revisionary from the perspective of the theoretical and linguistic practice at the time, even if the proposal correctly characterizes the subject matter.

A second such property is that of analyticity. I take analyticity to be a property of language. As such, a statement will be analytic not *per se*, but only relative to a linguistic practice at a time. For example, the sentence ‘Rape can only occur outside of marriage’ was plausibly analytic in 1800 but not in 2000. I agree with Haslanger, and for similar reasons to the ones she provides, that the sentence ‘Bachelors are unmarried men’, held up as the archetypical analytic statement in the philosophical classroom, is no longer analytic.³⁵ Quine was right when he said that no statement is immune to revision.³⁶ However, the dual-aspect theory of representation I have been suggesting allows us to distinguish analytic statements from conceptual truths in a way in which theories that identify linguistic meanings and concepts cannot. The status of a statement as analytic may be relative to a linguistic practice at a time, but conceptual truths are not; conceptual truths are eternal truths. The statements we accept as analytic, I suggest, are the statements that we take to express conceptual truths, and, of course, at the heart of a revisionary analysis is the claim that we erred in what we took the conceptual truths to be.³⁷

The distinction between linguistic meanings and concepts also has implications for the practice both of reporting what is said from across a theoretical divide and of reporting what is thought from across a theoretical divide. The case of reporting what is thought is relatively straightforward given my account. For example, Tilly, in 2000, can utter the sentence ‘In 1800 Abe believed that rape is not possible within marriage’, thereby attributing to Abe in 1800 the belief that rape is not possible within marriage. The practice of ascribing propositional attitudes to others is secured if the ascriber and the ascribee have access to the same thought, which, on the view I have proposed, is not jeopardized by a difference in linguistic meaning. The case of reporting what is said is *prima facie* more complicated, precisely because indirect speech reports are reports of what was *said* rather than of what was *thought*. As such, they appear to concern linguistic meanings rather than concepts, where linguistic meanings, I have argued, are precisely what differ across theoretical divides. How, then, can Tilly in 2000 report what Abe in 1800 said when he uttered the sentence ‘Rape cannot occur in marriage’?³⁸ Here we return to Frege’s initial claim that sentences express thoughts. Given this, we can say that Tilly’s report of what Abe said is true in virtue of the fact that it reports Abe as having uttered a sentence that expressed the thought that rape cannot occur in marriage. And this, we are assuming, is true.

³⁵ See the discussion in Haslanger (2006: section VII).

³⁶ See Quine (1951).

³⁷ The theoretical benefits of the distinction between analytic statements and conceptual truths will have to be explored in more detail on a future occasion.

³⁸ Saul (2006: 141) raises a version of this objection against a contextualist interpretation of Haslanger (2000).

8. Conclusion

In this chapter, I have urged the merits of a dual-aspect theory of representation that distinguishes the linguistic meaning of a term from the concept expressed by that term. Linguistic meaning, I suggest, is determined by patterns of actual use and tracks a community's understanding of a subject matter. Concepts, in contrast, are determined by real relations to objective properties and secure a subject matter about which there might be different understandings and different proposed analyses. I have argued that the distinction brings theoretical gains in a number of related areas. There is more to be said, but if I am right about the theoretical gains, then the view is worth taking seriously.

References

- Appiah, K. A. 1992. *In My Father's House: Africa in the Philosophy of Culture*. Oxford: Oxford University Press.
- Burge, Tyler. 1973. Reference and Proper Names. *Journal of Philosophy* 70 (14):425–39.
- Burge, Tyler. 1986. Intellectual Norms and Foundations of Mind. *Journal of Philosophy* 83 (12):697–720.
- Burge, Tyler. 1989. Wherein is Language Social? In Alexander George (ed.), *Reflections on Chomsky* (pp. 175–92). Oxford: Blackwell.
- Burgess, Alexis, and Plunkett, David. 2013a. Conceptual Ethics I. *Philosophy Compass* 8 (12):1091–101.
- Burgess, Alexis, and Plunkett, David. 2013b. Conceptual Ethics II. *Philosophy Compass* 8 (12):1102–11.
- Cappelen, Herman. 2018. *Fixing Language: An Essay on Conceptual Engineering*. Oxford: Oxford University Press.
- Carnap, Rudolf. 1947. *Meaning and Necessity. Meaning and Necessity: A Study in Semantics and Modal Logic*. Chicago: University of Chicago Press.
- Chalmers, David. 2003. The Nature of Narrow Content. *Philosophical Issues* 13:46–66.
- Chalmers, David. 2011. Verbal Disputes. *Philosophical Review* 120 (4):515–66.
- Clark, Andy, and Chalmers, David. 1998. The Extended Mind. *Analysis* 58 (1):7–19.
- Crane, Tim. 1991. All the Difference in the World. *Philosophical Quarterly* 41 (162):1–25.
- Fodor, Jerry. 1980. Methodological Solipsism Considered as a Research Strategy in Cognitive Psychology. *Behavioral and Brain Sciences* 3:63–73.
- Fodor, Jerry. 1987. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: MIT Press.
- Frege, Gottlob. 1884. *Die Grundlagen der Arithmetik: Eine logisch-mathematische Untersuchung über den Begriff der Zahl*. Breslau: Verlage Wilhelm Koebner. Trans. 1950. J. L. Austin as *The Foundations of Arithmetic: A Logico-mathematical Enquiry into the Concept of Number*. Oxford: Basil Blackwell.
- Frege, Gottlob. 1892. On Sense and Reference. Trans. 1948. In *Philosophical Review* 57 (3):209–30.
- Frege, Gottlob. 1918. The Thought: A Logical Inquiry. Trans. 1956. In *Mind* 65 (1):289–311.
- Gallie, W. B. 1956. Essentially Contested Concepts. *Proceedings of the Aristotelian Society* 56:167–98.
- Hasaday, Jill. 2000. Contest and Consent: A Legal History of Marital Rape. *California Law Review* 88 (5):1373–505.

- Haslanger, Sally. 2000. Gender and Race: (What) Are They? (What) Do We Want Them to Be? *Noûs* 34 (1):31–55. Reprinted in Haslanger (2012).
- Haslanger, Sally. 2006. What Good Are Our Intuitions? *Aristotelian Society Supplementary Volume* 80 (1):89–118.
- Haslanger, Sally. 2012. *Resisting Reality: Social Construction and Social Critique*. Oxford: Oxford University Press.
- Kaplan, David. 1989. Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives. In J. Almog and J. Perry (eds.), *Themes from Kaplan*. Oxford: Oxford University Press.
- Kuhn, Thomas. 1962. *The Structure of Scientific Revolutions*. Chicago: Chicago University Press.
- Orwell, George. 1949. *Nineteen Eighty-Four*. London: Secker & Warburg.
- Plunkett, David, and Sundell, Timothy. 2013. Disagreement and the Semantics of Normative and Evaluative Terms. *Philosophers' Imprint* 13 (23):1–37.
- Putnam, Hilary. 1973. Meaning and Reference. *Journal of Philosophy* 70:699–711.
- Putnam, Hilary. 1975. The Meaning of "Meaning". *Minnesota Studies in the Philosophy of Science* 7:131–93.
- Quine, W. V. O. 1951. Two Dogmas of Empiricism. *Philosophical Review* 60:20–43.
- Sainsbury, Mark. 2014. Fishy Business. *Analysis* 74:3–5.
- Saul, Jennifer. 2006. Gender and Race. *Aristotelian Society Supplementary Volume* 80 (1):119–43.
- Sawyer, Sarah. 2007. There Is No Viable Notion of Narrow Content. In B. McLaughlin and J. Cohen (eds.), *Contemporary Debates in the Philosophy of Mind* (pp. 20–34). Oxford: Blackwell.
- Sawyer, Sarah. 2010. The Modified Predicate Theory of Proper Names. In S. Sawyer (ed.), *New Waves in Philosophy of Language* (pp. 206–25). New York: Palgrave MacMillan Press.
- Sawyer, Sarah. 2018. The Importance of Concepts. *Proceedings of the Aristotelian Society* 118 (2):127–47.
- Segal, Gabriel. 2000. *A Slim Book about Narrow Content*. Cambridge, MA: MIT Press.
- Stalnaker, Robert. 1990. Narrow Content. In C. A. Anderson and J. Owens (eds.), *Propositional Attitudes: The Role of Content in Logic, Language and Mind*. Stanford: CSLI Press.
- Strawson, Peter. 1963. Carnap's Views on Conceptual Systems versus Natural Languages in Analytic Philosophy. In P. A. Schilpp (ed.), *The Philosophy of Rudolf Carnap* (pp. 503–18). London: Open Court.
- Sundell, Timothy. 2011. Disagreements about Taste. *Philosophical Studies* 155 (2):267–88.
- Sundell, Timothy. 2012. Disagreement, Error, and an Alternative to Reference Magnetism. *Australasian Journal of Philosophy* 90 (4):743–59.