

How Not to Change the Subject

Sally Haslanger

1. Introduction

The project of conceptual engineering, conceptual ethics, or conceptual amelioration – I will assume for the purposes of this paper that they are all the same – is highly contested. Some have argued that conceptual engineering – or whatever we call it – is impossible (Cappelen 2018), others embrace it enthusiastically (Burgess and Plunkett 2013). Debates over the possibility of conceptual engineering are confusing, however, because parties to the discussion start with very different accounts of concepts, meaning, content, and background philosophical methodology. There are also important differences in what those engaged in conceptual intervention see themselves as aiming for: Is the goal to introduce new (theoretical) vocabulary or revise the meaning of a (technical) term? Is the goal to promote a widespread linguistic change from the armchair? Is it an effort to change linguistic practices in the context of, and through the efforts of, a social movement?

Recently, I've argued that there are two different ways in which we might think of amelioration on an externalist account of content (Haslanger forthcoming). One form of amelioration is epistemic, the other is semantic. Epistemic amelioration is important and fairly straightforward; but sometimes semantic amelioration seems to be needed.

In this paper I will review the distinction between epistemic amelioration and semantic amelioration, and argue that at least one form of semantic amelioration – amelioration on functional grounds – is a coherent and sometimes valuable project. However, the point of conceptual *amelioration* as opposed to conceptual *replacement* is somewhat obscure. I shall suggest that the functional role of certain social concepts can sometimes warrant social critics in conceiving our project as amelioration rather than replacement.

2. Externalism about Content¹

In this section I will lay out the basics of an externalist account of (coarse-grained) content that will provide a backdrop for my arguments. I choose to proceed with these externalist assumptions, first, because I find the approach plausible; and second because I think it is a useful exercise to consider how amelioration might work within an externalist account of this sort. I am not the first to do this (see Cappelen 2018), but I hope to provide a different approach that allows us to be more hopeful about the possibility of amelioration.

On the view of content I endorse, our utterances and our mental states do not have senses or concepts as their content (Stalnaker 1998). We express, believe, suppose, (etc.) propositions, and propositions should be understood in terms of informational content, i.e., “as truth conditions, propositions as functions from possible circumstances to truth values, or equivalently, as sets of possible situations.” (Stalnaker 1998, 343). There are no “core commitments” associated with words that cannot be overturned or negotiated. Although in some sense we represent the world – propositions are abstract entities that carry information

¹ This section draws on my paper “Going On: Not in the Same Way,” forthcoming.

and are, to that extent, representational – the “mode of representation” is not part of the informational content of what we say and think. This allows you and me to think the same thing, the same proposition, even if we access what we are thinking differently.

How should we represent concepts on a view of this sort? Let’s suppose that the content of the concept is a partition of logical space.² We have access to partitions of logical space through different modalities, from different vantage points, and at different levels of granularity. So from a psychological point of view, *possession* of the concept may occur by virtue of different cognitive mechanisms and give rise to very different dispositions in different individuals. What it means to have a concept of X is not just what you can articulate, but how you respond to and coordinate with others in your environment, i.e., how your capacities for attention, categorization, interpretation, memory, language, inference, affect, and the like, are marshalled for the purpose to coordinating (and refusing to coordinate) with others in response to particular kinds of information. In short, to have a concept of X is to have a sophisticated set of capacities that enables one to process certain kinds of information about Xs. Consider Yalcin³:

To possess a concept is to have an ability to cut logical space in a certain way, to distinguish possibilities in terms of the sorts of things that answer to the concept....A concept determines a matrix of distinctions....To possess a concept, on this idea, is to be capable of entering states of mind sensitive to the associated distinctions. (Yalcin 2016, 14; also Pérez Carballo 2016, 466ff)

Yalcin’s model takes the relevant partitions to be partitions of worlds: “For example, the concept/subject matter BACHELOR corresponds to the partition of logical space distinguishing possibilities depending on what’s happening with the bachelors at each world— so that two worlds will belong to the same cell just in case they don’t differ in their bachelor respects.” (14) For our purposes, however, we can work with the more intuitive idea that the partitions serving as the content of concepts are of possible individuals rather than worlds. So the content of a concept (DOG) will be a function from a world (w_1, w_2, w_3) to a set of individuals in that world ($\text{dogs}_1, \text{dogs}_2, \text{dogs}_3$). Those who have the concept of DOG differ, however, in what features of dogs are accessible to them.

For example, I am a dog owner and love dogs. I can distinguish dogs from other animals. So can my dog’s vet. We both have the concept of DOG; our concepts have the same extension in each possible world. But the vet has a grasp of the logical partition of dog-space – the worlds of dogs – at a much finer granularity than I do, and so can answer much more detailed questions about dogs than I can. She is also likely to have a set of distinctions ready to hand that I lack entirely. The informational content of DOG can be grasped at different levels of resolution in a way that is responsive to the kind of questions that arise in employing it; of course, the vet’s questions are very different from mine. And, in turn, my questions are different from my neighbor’s, who also has the concept of DOG, but finds dogs annoying. His grasp of the

² Many people have developed this view in different ways. I think there is enough of a shared background so that one need not be a thoroughgoing externalist to accept much of what I say here. I adopt a Stalnakerian framework, but there are other ways of making the same points. Thanks to David Plunkett for pointing this out. Note also that Yalcin and Pérez Carballo (2016) are expressivists about certain kinds of content. I don’t, here, mean to embrace their full views but am simply drawing on the passages I cite.

³ In this paper, I’ll follow Yalcin in using SMALL CAPS for concepts. I’ll use *ITALICIZED SMALL CAPS* for partitions/subject matters which are the content of concepts.

concept allow him to distinguish “yippy” dogs from “scary” dogs – being driven by affect – but not much more.

In some cases we introduce different words to pick out the same partition of logical space for different purposes, e.g., vets might use the term ‘canis familiaris’ in certain contexts when their attention is focused on the biological features of dogs as opposed to wolves or other mammals. The informational content (partition of logical space) of CANIS FAMILIARIS and of DOG are the same, but the terminology reflects different contexts of inquiry. Yalcin makes this point with respect to the distinction between WATER and H₂O:

For example, the former [WATER] concept might (inter alia) be understood as embedding the partition into a subject matter reflecting parochial human interests and concerns, one including, say, the subject matter beverages— so that with the concept WATER, we (inter alia) locate that stuff amongst the beverages⁴ — while the latter [H₂O] embeds it in some part of the subject matter chemistry — so that with the concept H₂O, we inter alia locate that stuff amongst the chemicals. (15)

So here is the background picture: the *informational content* of a concept (and/or the meaning of a word) is a partition of logical space that divides possibilities.⁵ To *possess* a concept (and/or to grasp a meaning) is to have some cluster of capacities and mechanisms for using that grid of possibilities at some level of resolution, i.e., for making distinction(s), processing and storing the relevant information, answering questions. A concept, on this view, is not part of a proposition, nor is it a *thing* in the head; it is a cluster of dispositions to be responsive to differences in a particular region of possible worlds.

We are now in a better position to distinguish some different forms of amelioration.

- *Epistemic amelioration*: we improve our understanding of the informational content of the concept.
 - *Refinement*: We use concepts without having a very solid grasp of them. We may not be able to apply the concept to some possibilities, and there may be gaps in our judgments about cases. So we refine our concept based on a broader or deeper knowledge of the phenomenon, e.g., by gaining empirical knowledge, undertaking conceptual genealogy, and/or insight into logical space at a more fine-grained resolution.
 - *Experiential access*: we improve our access the informational content, gaining more reliable or illuminating access by different modes of presentation,. Those who train dogs have a different relationship with them than ordinary dog-owners. Knowledge how to coordinate with dogs, and experiential knowledge that comes with this, may prompt attention and responsiveness to

⁴ In some ways this is an odd suggestion, because it is not plausible to count ice or steam as a beverage, but both are, on the intended interpretation, water. To make this more plausible, I’ll understand ‘beverage’ to mean ‘beverage, when in liquid form.’

⁵ There are different ways to spell this out. For example, it might be a distinction between different sets of possibilities (an intension or modal profile (Schroeters 2015, 441)), or a distinction between sets of propositions with respect to their subject matter (e.g., Yablo 2014). Lewis and Yalcin take subject matters to be *partitions* of logical space; Yablo takes them to be *divisions* (2014, 36), replacing the role of equivalence relations in partitions with similarity relations. For our purposes, it isn’t crucial which option we take.

different features, e.g., being a scent hound or sighthound, that are important for the tasks at hand. (Hearne)

- *Informational/semantic amelioration*: we change what partition of logical space serves as the content of the term or concept, i.e., we undertake to change our thought and talk to do better in tracking reality. Better how?
 - *Alethic*: we are improving the resources available to track truths, e.g., a biological account of race prevents us from tracking important truths about race. Making these truths articulable using a social constructionist account can unmask ideology; it can also shine a light on new (emancipatory) possibilities. Ordinary scientific research can also shift content, e.g., biological theory can prompt changes in the distinction between animals and other kingdoms, with the result that the extension of ‘animal’ shifts.
 - *Pragmatic*: what we track with our language and our concepts can make life easier by shifting terms of coordination, e.g., ‘lunch’ once picked out a light meal at any time of day or night. Now when we invite a friend for lunch, we convey, with our term, information about the time of the day when we might meet.
 - *Moral*: because what we mean can affect what we do and what there is, semantic amelioration can also be (broadly) *moral*, e.g., if the informational content of (legal) ‘marriage’ excludes same-sex couples, this is a moral wrong.

3. Challenges for Semantic Amelioration

There are two important challenges to the possibility of semantic amelioration that I will take up in this paper: first, the “content as essence” objection; second, the “impracticability” objection. I’ll sketch the challenges in this section and will then address them each in the sections that follow.

a. Content as Essence

What is essential to a concept? Under what conditions do concepts simply alter (while persisting), and under what conditions are they generated, destroyed, or replaced? Recall that a concept has a particular content, but also, the possession of a concept brings with it certain capacities and dispositions. The partition of logical space is not something that comes and goes; but our grasp of it might. Some possibilities cannot even be imagined from a particular point in time. For example, who could have imagined self-driving cars several hundred years ago?

Consider two partitions of logical space D and D^* . Our concept of DOG, say, has D as its content: a particular set of all and only possible dogs. Suppose an ameliorator suggests that, instead, the concept of DOG has D^* as its content, i.e., a different set that is a proper subset of D . It might seem tempting to say that D^* can serve as the content for a concept of DOG, but that would be a different concept from ours because *our concept* has D as its content, and concepts have their informational content essentially.

If content is essential to a concept, then semantic amelioration is impossible, and the best we could hope for is concept replacement. Suppose that with more information about the genus *canis familiaris*, we decide that the informational content of DOG should be D^* rather than D (perhaps we find that what we supposed

to be a breed of dog – a breed universally treated as a dog – is really not a member of the species *canis familiaris*.⁶ It would seem that then the term ‘dog’ could be used for this new partition D^* , but the term would have a new meaning, and individuals who adjusted to the change would have a new concept. Concept replacement is compatible with retaining the word for the new content, so replacement could be masked. But the question is how concept amelioration is possible: can the content of a concept change, while it remains the same concept? Must we accept *content essentialism* for concepts?

b. Impracticability

Semantics provides an account of content; metasemantics provides an account of how a term or concept comes to have the content it has. The question: ‘What does ‘dog’ mean?’ is a question in semantics. The question: ‘By virtue of what does ‘dog’ mean dog?’ is a question in metasemantics. It would seem that in order to engage in intentional conceptual engineering, we would have to draw on a metasemantic theory about what determines content to decide what to fiddle with to get the change we are aiming for.

Cappelen (2018) argues that such intervention is impractical because of both “an epistemic point—that the metasemantics of our natural language terms are inscrutable—and a metaphysical point—that we have no control over the metasemantics.” (73) He continues:

Even if we had all the information about the metasemantics of a term (about the use patterns, the histories, the sources of information, the interaction between the experts, etc.), it would appear to factors that are in large part out of our control. For example, past facts play a role in determining the meaning of terms, but we can’t change the past. (74)

So the argument from intractability says that it may be that the content of a concept can shift, but this is not something we can control. Cappelen (2018, 73-76) argues, however, that we will inevitably keep trying, because we can’t help but theorize about what we *should* mean or how we *should* carve up the world. But thinking doesn’t make it so.

4. Epistemic Amelioration within a Two-Dimensional Approach⁷

One reason for accepting an externalist account of concepts is that it is useful to allow that individuals who have different modes of presentation, or different appreciation of possibilities, are still thinking (and talking) about the same thing. The ordinary person’s and the chemist’s thought and talk about water have the same content, even if they have different dispositions with respect to that content, e.g., they may be in a position to answer different questions and draw different conclusions from their WATER-thoughts. What we “have in mind” is not, ironically, in the mind, but in the world; the world provides us with shared

⁶ I am assuming here that the term ‘dog’ didn’t latch on to the kind *canis familiaris* by “reference magnetism” all along, but that our “universal” use of the term ‘dog’ to include the special breed is a significant factor in determining what the term means. The relationship between worldly joints and use in determining content is controversial and I will discuss it in connection with impracticability.

⁷ Note that the term ‘two-dimensionalism’ refers to different theories. Some, such as Kaplan’s and Chalmers’, are theories of meaning. Stalnaker’s is not. For Stalnaker, one dimension is semantic, the other is metasemantic. For a helpful overview of the different approaches, see Schroeter 2017.

content. However, sometimes we are interested not just in what is being thought, but what thoughts are easily available and how our thinking is embedded with other thoughts.

Let's return to the example of WATER and H₂O. Yalcin suggests that the content of WATER is the same as the content of H₂O, however, the two terms embed the content in different subject matters⁸, e.g., BEVERAGES, and CHEMICALS, respectively. What might this mean? In speaking of beverages, we distinguish water, juice, coffee, tea, beer, and such. In speaking of chemicals, we distinguish water, salt, gold, diamond, oxygen, and such. In thinking of water as a beverage, the relevant features are those that distinguish it from other beverages; *mutatis mutandis* for thinking of it as a chemical. Which features are relevant can be understood as relative to a background question, e.g., What do you have to drink?⁹ The features relevant to distinguishing the options are then salient and accessible to the thinker.

Suppose that an agent is a question and answer machine. Its belief and knowledge are to be understood as capacities and dispositions to answer questions. Such a machine will need some mechanism or mechanisms for storing information, and some way of using that information to generate answers to the questions it receives. *How* it stores the information is not directly relevant to *what* it knows or believes. (Stalnaker 1991, 436, my italics)

Let's take content of the concept BEVERAGES to be a partition in logical space between BEVERAGES and NON-BEVERAGES. This partition has multiple cells, one of which is the water cell. CHEMICALS are a different partition, and one cell is also the water cell. In both cases, the water cell is the same. But we access the cell differently and attend to different features, depending on the context. To possess the concept WATER, one need not know anything about its chemical composition (as is obvious from the fact that Aristotle possessed the concept WATER but had no knowledge of modern chemistry). Plausibly, "...the accessibility of knowledge and belief can be understood only relative to the actions they are being used to guide" (Stalnaker 1991, 439). As a result, amelioration – typically epistemic amelioration – allows us to access features of the object of our belief asking different questions, situating our beliefs in relation to other beliefs, and undertaking different tasks.

Stalnaker (and others) have used a distinction between *what* we believe and *how* we believe to explain why, on an externalist account of content, we are not logically omniscient. They have also explained how we can fail to have knowledge of identity claims – such as Hesperus = Phosphorus – and how they can appear to be contingent. I will sketch this approach and extend it to thinking about concepts. I will then argue that once we see how the content of our thought adjusts in relation to worldly considerations, we can make sense of the possibility of semantic amelioration.

Suppose O'Leary does not know that Hesperus = Phosphorus. So we may want to tell him. But it is puzzling how our contribution is helpful. If 'Hesperus' and 'Phosphorus' both refer to Venus, and if O'Leary knows the meaning of 'Hesperus' and of 'Phosphorus,' then 'Hesperus = Phosphorus' expresses a necessary truth that we should expect O'Leary to know simply by knowing the meaning of the terms.

⁸ Note that Yalcin is drawing on a discussion of subject matter according to which subject matters are partitions of worlds, not individuals, so the suggestion of inclusion he makes here is not as straightforward if we are thinking of partitions of individuals rather than worlds. This shouldn't matter for our purposes.

⁹ Below, Stalnaker proposes an image of us as question and answer machines. It is important to see that even on Stalnaker's view, the machine's questions and answers are primarily action oriented.

However, drawing on the considerations outlined above, it is plausible to say that O’Leary’s access to Hesperus is different from his access to Phosphorus, and because of this he cannot access the identity.

Stalnaker situates the problem within a theory of conversation – considering the different presuppositions that serve as part of the background to each interlocutor’s contribution. For our purposes, we don’t need the full detail of his account, so I will focus on his two-dimensional approach to highlight the interaction between semantics and metasemantics. Consider an utterance U . Semantics gives us a theory of the content of U ; metasemantics gives us a theory of how U has the content P it does, i.e., by virtue of what U expresses P . As Stalnaker points out:

We need two dimensions since we start with the fact that the truth value of a proposition (at least a contingent proposition) depends on the facts. But since the identity of the proposition expressed in a given utterance also depends on the facts, the truth value of the utterance will depend on the facts in two different ways: first, the facts determine what is said; second, the facts determine whether what is said is true. (Stalnaker 2004, 302)

To capture the complexity, Stalnaker introduces the idea of a *propositional concept*, using a two-dimensional matrix. Along the vertical axis we list the circumstances relevant to determining content; along the horizontal axis we list the circumstances of evaluation. In the case of propositions, the circumstances are represented as worlds. Let w_1 be our world, where ‘Hesperus’ and ‘Phosphorus’ both refer to Venus. Let w_2 be the world O’Leary thinks it is, where ‘Hesperus’ refers to Venus and ‘Phosphorus’ refers to a different heavenly body, say, Mars. We can then represent the options for the utterance ‘Hesperus = Phosphorus’ this way:

	w_1	w_2
w_1	T	T
w_2	F	F

In other words, if ‘Hesperus’ and ‘Phosphorus’ mean what they mean in our world, then ‘Hesperus = Phosphorus’ expresses a necessary truth. If ‘Hesperus’ and ‘Phosphorus’ means what they mean in a world like O’Leary thinks ours is (w_2), then ‘Hesperus = Phosphorus’ expresses a necessary falsehood. When we tell O’Leary that ‘Hesperus = Phosphorus,’ we communicate to him that we are in w_1 rather than w_2 . Stalnaker captures this by saying that the information O’Leary lacks is best captured by the diagonal proposition: the proposition that is true in our world, w_1 , but would be false in w_2 because we hold fixed *our* referents, as determined in w_1 but consider the evaluation in a world where two different heavenly bodies occupy the positions from which O’Leary views them. As a result, the utterance ‘Hesperus = Phosphorus’ can appear contingent, because facts that determine its content could have been different.

The direct reference theory of names gave an externalist account of the facts that determine reference: statements containing names have the content that they have because of the way speakers using them are causally connected with things in the world. Consequently, in possible worlds where the astronomical facts are different, the semantic values of names referring to astronomical bodies may be different, and so different propositions may be expressed with those names. (Stalnaker 2004, 302)

Our concern is not propositions, however, but concepts. To extend the model, let's run through this first, switching from propositions to names, and then moving from the content of names to the content of concepts. Consider O'Leary again. He thinks that the content of 'Phosphorus' is Mars. Again, this may seem odd. If the content of 'Phosphorus' is Venus, then does he believe that Mars = Venus? Surely not. In thinking that Phosphorus is Mars, we don't need to attribute to him an incoherent belief. Instead we attribute to him a belief based on a different supposition about how things are in our world, in particular, about how the content of 'Phosphorus' is fixed. As before, the worlds along the vertical axis represent the metasemantical facts that determine the referent of the term; the worlds along the horizontal axis represent the context of evaluation. His understanding of the content of 'Phosphorus' can be represented in this way:

	w_1	w_2
w_1	Venus	Venus
w_2	Mars	Mars

O'Leary is plausibly aware that if 'Phosphorus' picks out Venus, then 'Phosphorus = Venus' is a necessary truth and its denial is a necessary falsehood. But he imagines that 'Phosphorus' picks out Mars, i.e., that he is in w_2 . The appearance that Phosphorus could have been Mars is due to metasemantical contingency, viz., the diagonal idea that although we use 'Phosphorus' to pick out Venus (w_1), we could have used it to pick out Mars (w_2).

With this model in place, the move to concepts is relatively straightforward. Suppose Asha visits the United States regularly. In the spring she usually visits Massachusetts, and in the summer she visits Kentucky. In her visits to Massachusetts she comes to love the redbuds¹⁰: their bright pink flowers poking directly out from the bark, their broad open shape, their relatively smooth bark. In Kentucky, she loves to sit in a park under a spicewood tree with broad heart-shaped leaves with a distinctive reddish hue that provide a perfect canopy against the hot sun. She has no idea that the spicewood and the redbud are the same species: *cercis canadensis*. One afternoon in the summer, she and a friend decide to stop for an iced tea; they choose a table under a tree, and the friend comments on the wonderful shade of the redbud. Asha looks up and sees a spicewood. She then puts it together: redbuds are spicewoods!

What was she thinking? If the content of 'redbud' and the content of 'spicewood' are both *cercis canadensis*, then, before teatime, was she thinking that the species is different from itself? That would be incoherent! But as in the 'Hesperus = Phosphorus' case, we can imagine a possible scenario in which the tree species that has such a distinctive form of flowering (what Asha thinks of as a redbud) is not the tree species that has such a beautiful leaf canopy (what Asha thinks of as a spicewood). This is the world that Asha thought she was in. The problem, as before, is that if redbuds and spicewoods are both *cercis canadensis*, then no such world is possible. But also, as before, we can use the two-dimensional approach to model this.

Think of tree species as partitions of possible individuals: possible trees of the species in question. The partition *cercis canadensis* distinguishes trees in that species from everything else. In our world – w_1 – both 'redbud' and 'spicewood' have their reference fixed to the same partition: *cercis canadensis*, i.e., the content of both is that partition. However, we can imagine another world – w_2 – where 'redbud' is used for *cercis*

¹⁰ I assume that she loves the Eastern Redbud, i.e., *cercis canadensis*. There are, by some counts, twelve species of redbud (*cercis*), but the one that she would see in Virginia and is commonly called 'redbud' there is the Eastern Redbud.

canadensis ('cc'), and 'spicewood' is used for a different species that looks like redbuds in the summer, call it *cercis kentuckiense* ('ck').¹¹ This world differs from ours in the metasemantical facts. Not knowing that spicewoods are redbuds, Asha seems to think that the content of 'spicewood' is *cercis kentuckiense*, i.e., that she is in w_2 .

	w_1	w_2
w_1	cc	cc
w_2	ck	ck

When we characterize what Asha was thinking before she learned that spicewoods are redbuds, we don't have to say that she was thinking, incoherently, that the content of the two terms, 'redbud' and 'spicewood,' viz., the *cercis canadensis* partition, is different from itself. Instead we can say that she was thinking that 'spicewood' had as its content a species other than *cercis canadensis*, perhaps *cercis kentuckiense*, i.e., that she was in world w_2 . The appearance that spicewoods could have been *cercis kentuckiense* is due to metasemantical contingency, i.e., the diagonal idea that although we use 'spicewood' to pick out *cercis canadensis* (w_1), we could have used it to pick out *cercis kentuckiense* (w_2).¹² Returning from concepts to propositions, we might say that her mistaken belief was the vertical belief that *cercis canadensis* = *cercis kentuckiense*.

How plausible is this? Drawing on our earlier discussion of Yalcin, the idea is that we access a partition of possible dogs, or trees, or beverages, at a certain level of granularity. The partition of possible dogs may have two cells (yippy/scary) or 190 (the official AKC breeds), or any number of other options. Asha's interaction with trees prompted her to consider them seasonally: Which are beautiful in the spring? Which are excellent sources of shade? Redbuds were a cell in the former. They were also a cell in the latter. But she had different ways of accessing the cell, depending on the context. In learning that redbuds are *cercis canadensis* she learned that the cells in each way of dividing the partition are the same cell. By virtue of her conversation under the shade of the tree, she also learned that there was something incomplete or inapt about her understanding of redbuds.

That different questions lead us to carve the world differently, and that action often requires us to resolve the different carvings, is an important source of knowledge.

It seems to be a fact, not a problem, that questions - even nonleading questions without presuppositions - can change what we know and believe by bringing out what was previously merely implicit in what we believed. (Stalnaker 1991, 438)

5. Semantic Amelioration

In the last section I suggested a way to think about epistemic amelioration within a two-dimensionalist approach (using a metasemantic interpretation). The task before us now is to determine whether it is possible for a concept to shift in its content – for it to have a different partition of logical space as its

¹¹ *Cercis kentuckiense* is a made-up (possible) species. More realistically, she might mistake the spicewood is a kind of catalpa, e.g., catalpa x erubescens 'purpurea'.

¹² I am interested in connecting this two-dimensional approach to Epstein's (2015) distinction between grounds and anchors), but don't have the time or space to do it here.

content – while remaining the same concept. In other words, on an externalist view of concepts (of the sort I've been working with), must we accept *content essentialism*? In this section, I will argue that there are circumstances when what we learn about the world pushes us to rethink and revise the content of a concept, without replacing it. Very broadly, not only our grasp of semantic content, but also our metasemantics can be corrected by the world.

Let's begin with an example discussed by Stephen Yablo. The point of the example, for our purposes, is to highlight how meaning depends on the facts, as Stalnaker suggests, in two different ways: "first, the facts determine what is said; second, the facts determine whether what is said is true." We can be wrong about the facts in both ways. As a result, we are faced with choices when we discover the facts.

Suppose we are fixing the referent of 'meter.' Kripke suggests that we can do so by way of an accidental property of a particular stick (Kripke 1980, 55). Yablo argues:

Since it is an empirical matter whether stick S is "the length he wants to mark out", we need to ask what happens if he is wrong and it is a different length than intended. It might be, for instance, that the stick is a millionth of an inch long, but emitting magnification rays that delude us into seeing it as longer. Or maybe the stick is a mile long, but much farther away than anyone had realized. I take it that it is no part of the reference-fixer's understanding of 'meter' that it continues to stand for the length of S even if S is much shorter or longer than it appears. Since this cannot be a priori ruled out, we don't know a priori that the stick is a meter long if it has a length at all. (Yablo 2008, 184)

In Yablo's example, the baptism of the meter is unsuccessful because crucial background presuppositions are unsatisfied. Suppose instead that we succeed in fixing the referent of 'meter' using the stick – let's say the stick that happens to be exactly 39 inches long. We find, though, that using a particular physical object as the standard is a nuisance because it doesn't achieve the uniformity of measure that is required – not everyone has access to that particular stick. It is reasonable, then, to correct ourselves and at the same time correct the content of 'meter.' We might find, for example, that there is a distinctive physical phenomenon that is close in length (say, 39.3701 inches), and change the length of a meter to that. In changing the length, are we abandoning the concept of METER and introducing a new concept METER*? Plausibly not. The initial length fixed didn't work for us and the second did.

One might complain that this example is too far-fetched. But in fact, the baptism of the meter has happened at least six times since 1798, using different devices to fix the referent to make the unit of measure more stable, precise, and useful (Cardarelli 2003, cited in Wikipedia "History of the Metre"). Although the adjustments of the term 'meter' were, in some cases, tiny, they did change the relevant partition of logical space that is the content of 'meter.' This shows, I submit, that we can meaningfully claim that the concept of METER changed its content over time, i.e., the informational content changed but the concept was not replaced with a new concept. This is not to say that we *must* say that there has been one concept of METER through these changes. That will depend on what questions you are asking and what subject matter is your concern.

This example shows that there is a way of thinking about concepts that allows them to alter their content without being replaced. But there are limitations to the example. Most importantly, it is a case of stipulated content where the intention behind the stipulation is known, i.e., there is an agreed upon target

and a specific use for that target. Concepts of the sort we are mostly interested in ameliorating are not stipulated. Yet the phenomenon is more natural than one might think from the example. Philosophy of language often works with simplified examples where meaning is clear in order to work out a model. But meaning is not always clear and the need to take into account past understandings, newly discovered facts, and pressures on ongoing coordination, require complicated judgments that don't necessarily fit a single pattern.

The questions that guide the ameliorator are: Is this content/partition what we *should* be thinking and talking about? Is this what we *need to* think and talk about to carry on meaningfully, rationally, morally in order to answer the pressing questions? Recall the example of WATER and H₂O. If the concept WATER situates and compartmentalizes our WATER-thoughts in relation to beverages, and if H₂O situates and compartmentalizes our H₂O-thoughts in relation to chemicals, then the aptness of the division for those uses may, over time, lead to shifts in content. For example, beverages should be drinkable; not all H₂O is drinkable (some is ice, some is steam, some is contaminated, some is in micro-droplets). Even now I think there are contexts in which it can sound odd to say that steam or ice is water. (“Quick, she needs water!” (Said of a dehydrated person.) Would offering H₂O in any form – blasting her with steam? – be an apt response?) As climate change progresses and the value of potable liquid water in sufficient quantities to drink to remain hydrated becomes a pressing need, it would not be surprising if the term ‘water,’ at least in some contexts, shifted its content to a subset of the H₂O partition. Would we have changed the concept WATER? If it is more important to the individuation of the concept that it be responsive to the subject matter in which it is situated, than that it maintain its exact content, then we should say that a change of content is compatible with the continuity of concept. In other words, if concepts are not to be understood simply as cognitive placeholders for content, but, e.g., provide frames for content at different levels of granularity, highlight features of the content in response to questions, and situate the content in different subject matters, then these aspects of concepts should carry weight in individuation. The whole point of having concepts in the model is to do justice to the different ways we access content and our different dispositions to respond to content – differences across individuals, differences for individuals at different times, and differences in possible scenarios; to ignore these differences in considering the individuation of concepts would be to sacrifice what we have gained.

6. Functions

I have suggested elsewhere (Haslanger 2012, Ch. 6) that we should look to the point, or purpose, or function of a concept to answer the motivating questions mentioned above: Why do we need this concept? What function does it serve? Should *this* be what we think and talk about, given our legitimate purposes? If a concept has a particular function, and the content associated with it fails to carve the world in a way that enables the concept to fulfill its function aptly, then it would be reasonable, I hope, to change the content. In the case of water and H₂O, we considered the idea that the concepts have different functions in our cognitive economy, so modification of the concept of WATER might be warranted when modification of the concept of H₂O is not. Herman Cappelen (2018, Ch. 16) has argued, however, that there isn't a meaningful sense of function that would serve this purpose. This is also relevant to the limitations of the meter example above, for in that case we could talk about intentions and purposes, but in most cases of language use, there are no moments of baptism with clear targets, no human designer deciding what words to invent.

Amie Thomasson (forthcoming) has made the important point in response to Cappelen that there are conceptions of function that do not depend on human intentions or purposes. It is clear, for example, in the biological study of ecosystems that certain species have a function relative to the ecosystem. Thomasson considers this with respect to language, suggesting a Millikan-style interpretation of function. She suggests that as a first step, we should aim to:

...identify something that this range of concepts *does* or (better) *enables us to do*, that we couldn't do (or couldn't do as effectively or efficiently) without it (or an apt translation). Analyses like these can serve as clues to proper function analyses: to why it would have been useful to have concepts like this, why terms that express them might have been perpetuated in our culture. (12)

This starting point does not assume that we are the designers of the concepts. Cappelen (2018) responds that:

The reason 'salmon' is useful for us is that it can be used to talk about salmons (or denote salmons). The reason 'freedom' is useful is that it can be used to talk about freedom. We care about salmons and freedom and so we have words that enable us to talk about them. (Of course, all of this could have been done by other words—so we don't need any of those terms to perform that function.) However—and this is the key point—beyond these disquotationally specified functions, there's variability. We can use 'freedom' in speech acts that have as their aim to undermine freedom or promote it or discuss it or disparage or make fun of it or . . . There's no limit to what we can go on to do with this term. These activities will vary wildly between contexts and over time. If the goal is to find functions that are more substantive and informative than the disquotationally specified functions, then it will be unsuccessful. (187)

Cappelen seems to be making two points. First, the primary function of terms (or, given that we are talking about conceptual amelioration, concepts) is to latch onto something in the world so we can talk/think about it. And what the words latch onto can only be captured disquotationally. Second, bits of language and thought many have many different functions, and there is no way to specify which function we should “do justice to” as we ameliorate other than latching onto the relevant bit of the world.

Thomasson takes exception to both of these points. For different reasons and against the backdrop of different semantical views, she and I agree that concepts do much more than latch onto partitions of logical space. Note, for example, that in the social domain, what is there to be “latched onto” may not exist independent of the resources we use to “latch.” What symbolic resources we have, and what there is, are much more closely related, and the direction of fit may reverse. So then we must evaluate our language, not only by whether it properly “latches” but what in the world it makes possible. Moreover, it is not at all clear why Cappelen would think that disquotation is the only way to capture what we are picking out with a term. One of the great features of our language is that we can construct phrases to pick out things and kinds in many different ways.

More relevant to our purposes here, however, is whether ordinary concepts can have a function by reference to which they should be evaluated. Cappelen's worry is that any concept has a large and indefinite number of functions, and there isn't a way to identify *the function* by reference to which such an evaluation makes sense. As Thomasson points out, Cappelen seems to think that if an object functions in

a particular way for a particular purpose, even on a single occasion, then we can say that the object has that function. So if a screwdriver is used to open a paint can, it has the function of can opener; if it is used to prop open a door, it has the function of a door stop, etc. Drawing on her work on artifacts, Thomasson suggests that we should not consider just any use of something as a basis for attributing “proper function”. In the context of language, a Millikan-like etiological approach might provide a way to distinguish the myriad of possible uses of a term or concept from its proper function.

I am less inclined to reach for proper functions in understanding the function of concepts. Rather, I prefer to adopt a systems approach to understand the phenomenon.¹³ On this approach, something has a function relative to a system: to attribute a function to it is to identify its role in producing selected features of the system in question. The paradigm systems are typically organisms, ecosystems, and machines. A heart functions *in a living animal* to pump blood; a carburetor functions *in an internal combustion engine* to mix air and gasoline; bees function *in corn agriculture* to pollinate the corn. Relative to the system of corn production the bees are pollinators; relative to the system of honey production, they are the producers. Of course the two agricultural systems (corn and honey) are interdependent, but our questions may focus attention on one function, relative to one system or another. It is crucial to note that what is being explained is the capacities of the system, e.g., that yield corn, or honey, in terms of the capacities of the system’s parts. Cummins (1993) calls this form of explanation “functional analysis.” Attributing functions in a system analysis does not commit us to proper functions in an etiological sense (Godfrey-Smith 1993, 1994).

On a systems approach, not all parts of a system serve to sustain it; some parts may destabilize it and other parts counter the tendency to destabilization. The process of restabilization may not return the system to the original set point, but allow it to evolve to a new set point. Most systems manifest dynamic rather than static homeostasis, i.e., they are dynamical systems. The choice of system by reference to which we identify a function is *our* choice: we are the ones asking the questions; we are the ones seeking explanation of regularities that matter to us. The fact that the heart functions to pump blood relative to the circulatory system is an objective fact, but whether we attend to this function rather than other functions the heart may have in other systems is a result of our interest in explaining and repairing the circulatory system. Recall the water/H₂O example: we might situate the WATER concept within a beverage system and the H₂O concept within a chemicals system. Then the concepts function differently, relative to those systems, even though they have the same content. The concepts reflect our capacity to be responsive to different partitions in different ways, and responsiveness based on different modes of access, background knowledge, and affective engagement, give rise to different actions.

The regularities that matter to me are social regularities. Social theory often pursues a systems level of analysis. Why does the social system as a whole have certain properties, e.g., why (and how) does the system create and/or sustain gross disparities of wealth, or certain patterns of domestic violence? The explanatory project is to identify relevant parts of the system and show how their interactions function to have the result

¹³ There is an extensive literature on functions. Broadly, the accounts fall into purposive, etiological (Wright 1973, Millikan 1989), and systems (Cummins 1975) approaches.

in question.¹⁴ Note that how the parts function may not be how people think they function. In fact, this is often how ideology works: we participate willingly in systems whose effects we abhor because we don't see how the practices we engage in contribute to the pernicious features of the system.

Currently social, legal, and religious systems lay claim to the concept of MARRIAGE. But the concept has different functions relative to those systems. For generations, there were assumptions about the concordance of the systems. In some historical contexts, a social marriage, was a legal marriage, was a religious marriage (more or less). Legal and religious marriage can be stipulated, in a sense, and social marriage was normatively constrained to conform to the requirements of legal and/or religious marriage. But over time, these systems have, in practice, fallen increasingly out of concordance. For example, in some religious communities, legal marriage was once both necessary and sufficient to be married within the church, but now legal same-sex marriage does not count as “real” marriage anymore. One issue, then, is whether one path through these systems counts as the continuity of a single previous concept, and the others not.

There seem to be at least three options. (i) The different forms of marriage were assumed to be concordant, but weren't. The concepts always had different content. (ii) The different forms of marriage are still concordant – the concepts have the same content – but certain religious traditions (or legal jurisdictions) fail to understand what marriage really is. (iii) The concepts provided different perspectives on the same content – relative to the different social/legal/religious systems – and over time the different pressures on the systems have led to changes in the content relative to one system but not others. (See also Khoo 2018.)

Our focus is on option (iii), for it suggests a semantic shift. For centuries in the West, marriage was assumed to be formed through a religious/legal contract between “one man and one woman.”¹⁵ Within this historical tradition, marriage regulates the possibilities of sexual intimacy; it creates a private space for the creation of and care for children; it produces “basic” economic units. It also interacts with other sub-systems of food production, education, healthcare, etc. Suppose we are interested in a social system, and want to know why, say, job segregation by sex/gender among Whites was so entrenched. Drawing on Goldin (1980), consider the period between 1870-1922 in the United States. Goldin argues that White working women in that period were primarily unmarried, semi-skilled, and anticipated leaving the job market when they married. As a result, there were fewer opportunities or motivations for women to develop skills that would bring higher wages and sustain them through a lifetime; jobs that allowed this

¹⁴ I find it plausible to think of such explanations as pointing to mechanisms that give rise to the features of the broader system. Tilly (2001), however, seems to contrast “mechanistic” explanations and systems explanations. (See also Tilly and Tarrow 2015, Ch 2.) My reading of Tilly's resistance to systems explanations is that he assumes that such explanations involve an etiological approach to function: “...systemic explanations, strictly speaking, consist of specifying a place for some event, structure, or process *within a larger self-maintaining set of interdependent elements* and showing how the event, structure, or process in question serves and/or results from interactions among the larger set of elements” (2001, 23, my emphasis). Thanks to David Hills for reminding me that not all social practices, or social structures, are self-maintaining. Some, such as resistance movements, are even designed to be self-destructive: success will make the movement obsolete.

¹⁵ Alex Byrne has pointed out exceptions (including Caligula's marriage to his horse), but the exceptions need not undermine the claim that heterosexual marriage was the paradigm case that served to fix the referent.

were given to men. The fact that marriage was the primary institution for managing childcare, together with the gendered distribution of labor in marriage, helps explain the pattern of gendered job segregation. Marriage functions in that system to entrench job segregation. This is not to say that it is the sole cause of gendered job segregation; the analysis of the system points to it as one factor among others.

Systems can be evaluated on a variety of grounds, e.g., stability, endurance, efficiency. Social systems can also be evaluated on (broadly) moral grounds. One can argue that social systems built on the nuclear family, that distribute childcare labor almost entirely to women in the home, and prevent women from having their own income, are bad for women. In such cases women's options for exiting a marriage are limited and costly. This gives men in the relationship greater power. Expectations of power increase the probability of domestic violence. So we have reason to change the expectations and practices to disrupt such a system through social, political, and cultural interventions.

It is also the case that marriage as it was long practiced – between “one man and one woman” – also functioned to marginalize same-sex couples, deny them the rights granted to heterosexual couples, prevent LGBTQ partners from adopting children (leaving many children without family support), and had a detrimental impact on the health and well-being of many. Here too we have had reason to change the system. However, in this case, the concept of MARRIAGE seemed to be a barrier to change. Many people took the explicit requirements of legal and religious marriage to set requirements on social marriage, because of the congruence mentioned above: access to the partition of married couples was shaped by the significance of that partition in legal and religious institutions. Within queer communities, however, same-sex marriage began to happen prior to its legalization. Some of the social and interpersonal functions of marriage were being fulfilled by new kinds of couples, even if the legal and religious functions of marriage were not. As same-sex couples came to be increasingly recognized as socially married, pressure increased to extend the boundaries of legal and religious marriage.

We can make sense of this within a two-dimensional framework by noting pressure to accept vertical adjustments in the metasemantics. My goal here is not to argue that we should take the concept of MARRIAGE, specifically, to have changed its content, i.e., for the actual semantic amelioration of MARRIAGE. I am arguing, generally, for the *coherence* of semantic amelioration against the content essentialist, and using a story about MARRIAGE (perhaps true of our world, but maybe only of another possible world) as an example. It is coherent to claim that there is a contingency in the metasemantics of a word or term and that it can turn out that particular metasemantic events that fix the content at a given time or world gets things wrong (as in the case of ‘meter’). In the case of marriage, perhaps the metasemantics that fixed the referent of ‘marriage’ by reference to a particular heterosexual social formation succeeded in establishing a content for ‘marriage’ for a given period of time. But we have learned over time that focus on this phenomenon distorts our access to a broader and more important partition that includes same-sex couples, attention to which would allow us to better manage intimacy and childrearing. Given this, we should undertake (and have undertaken) metasemantic efforts to shift the content of our concept – the concept situated in the subject matter of *social* (and eventually legal and religious) relations – to this more inclusive partition. After the change of content, MARRIAGE counts as the same concept because it situates the new partition of couples in the social system where the old partition was situated.

7. *Impracticability*

On Cappelen's (2018) view, metasemantics is a mysterious business. For most of our terms, reference was determined long ago, shrouded in the mists of history. The process is "incomprehensible and inscrutable" (53), and is not something we can change anyway. Conceptual engineering is not even, strictly speaking, a project that makes sense on his account, for "The kind of thing philosophers and psychologists call 'concepts' plays no role in my theory." (53) No wonder he thinks that conceptual engineering is impracticable.

I've sketched an externalist view of content that makes room for concepts. Concepts are not things in the head, but they are capacities we have for accessing partitions of logical space at different levels of granularity for different purposes. Cognition is shaped by social practices and social structures (Zawidzki 2013). Socialization, language learning, and the like, can affect what partitions are available for thought and how we access their members. Education brings our cognitive capacities into alignment on certain ways of understanding what we are talking and thinking about so that we can understand, more or less, what others mean and predict their behavior.

Of course, proposals for meaning change offered in philosophy journals do not change the meanings of ordinary terms (though philosophers and other theorists can stipulate meanings of new terms). *Of course*, philosophers, doing what philosophers usually do, do not bring about social change. Social change is produced by social movements, by changes in law and policy, by capitalist investment, by climate change, by random contingencies in a dynamic social system. Philosophers, however, are sometimes engaged in social movements that promote changes in our practices. They are sometimes engaged in working with social movements to determine what changes would result in a society that is more just and less harmful and how these changes might be brought about. Conceptual engineering can be part of such politically engaged activity. In other words, there is room for philosophical work within social movements.

I haven't given here a full-blown metasemantics. As mentioned above, my first goal has been to illuminate how conceptual engineering is a coherent project. I've argued that in the case of conceptual tools that we design and implement, it has occurred (recall the 'meter' example). In the case of ordinary concepts that emerge more organically, I have argued that reference adjustment can occur when we learn that concepts (cognitive capacities) play a role in enabling us to organize our lives together, and that we can learn that the cognitive capacities we have developed are distorted, either in their content or in the mode of access to the phenomenon that we should be tracking. They are functioning to sustain a social system in a way that is problematic, so we have reason to change either the mode of access (epistemic amelioration) or the content (semantic amelioration).

In the process of change, it is often not obvious whether the best strategy is to introduce a new term for a new content, to appropriate the old term for a new content, or to shift the content of the term we have been using all along. I have not given clear criteria for determining when we should opt for one or another of these moves, or how to clearly distinguish which is being attempted. I have suggested, however, that there are times when ordinary terms have a tight association with a particular subject matter ('water' with beverages rather than chemicals). Because this tight association can be hard to break, it can be strategically useful. The fact that 'marriage' is tightly associated with a social, legal, and religious relationship allowed the recognition of social *marriages* in queer communities to push us toward legal and

religious recognition. If we had instead only recognized queer ‘life partnerships,’ then it is likely that the struggle for equal rights and respect would have, at best, been delayed. But how we proceed in any particular case will depend on the history, the politics, and the goals at issue.

I agree with Cappelen that we cannot control the process of change, but we can begin to accept new paradigms that re-fix the referent, we can incorporate new things into our practices in ways that alter what partitions we take to be relevant to carrying on. We, philosophers, don’t need to just sit in armchairs. We can be part of social movements by illuminating different parts of the social world, articulating normative demands on cognition, and giving people cognitive access to new possibilities.

Acknowledgements: Thanks for discussion on topics covered in this paper from: Alex Byrne, Herman Cappelen, Brian Epstein, Samia Hesni, Justin Khoo, Dan Lopez de Sa. Special thanks to Stephen Yablo for helpful feedback and suggestions.

References

- Burgess, Alexis and David Plunkett. (2013). “Conceptual Ethics” (I and II), *Philosophy Compass*. Vol. 8, No. 12, 1091-1011 and 1102-1110.
- Cardarelli, François. (2003). *Encyclopedia of Scientific Units, Weights, and Measures*. Trans. M. J. Shields. London: Springer-Verlag.
- Cappelen, Herman. (2018). *Fixing Language*. Oxford: Oxford University Press.
- Cummins, R. (1975). “Functional Analysis.” *Journal of Philosophy* 72: 741-765.
- Epstein, Brian. (2015). *The Ant Trap*. Oxford: Oxford University Press.
- Godfrey-Smith, Peter. (1993). “Functions: Consensus Without Unity.” *Pacific Philosophical Quarterly* 74(3): 196-208.
- _____. (1994). “A Modern History Theory of Functions.” *Noûs* 28(3): 344-462.
- Goldin, Claudia. (1980). “The Work and Wages of Single Women, 1870-1920.” *The Journal of Economic History* 40(1): 81-88.
- Haslanger, Sally. (2012). *Resisting Reality: Social Construction and Social Critique*. Oxford: Oxford University Press.
- _____. (forthcoming). “Going On, Not in the Same Way.” In *Conceptual Ethics and Conceptual Engineering*, ed. A. Burgess, H. Cappelen, D. Plunkett. Oxford: Oxford University Press.
- Khoo, Justin. (2018). “Quasi Indexicals.” *Philosophy and Phenomenological Research*, <https://doi.org/10.1111/phpr.12519>
- Kripke, Saul. (1980) *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- Millikan, R. G. (1989). “In Defence of Proper Functions.” *Philosophy of Science* 56: 288-302.
- Pérez Carballo, Alejandro. (2016) “Structuring Logical Space.” *Philosophy and Phenomenological Research* 92(2): 460-491.

- Schroeter, Laura. (2017). "Two-Dimensional Semantics", *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2017/entries/two-dimensional-semantics/>>.
- Schroeter, Laura and Francois Schroeter. (2015) "Rationalizing Self-Interpretation." In Chris Daly, ed., *The Palgrave Handbook of Philosophical Methods*. Basingstoke: Palgrave-Macmillan, pp. 419-447.
- Stalnaker, Robert. (1991). "The Problem of Logical Omniscience, I" *Synthese* 89: 425-440.
- _____. (1998). What Might Nonconceptual Content Be? *Philosophical Issues* 9: 339– 352.
- _____. (2004). "Assertion Revisited: On the Interpretation of Two-Dimensional Modal Semantics" *Philosophical Studies* 118: 299–322.
- Tarrow, Sidney and Charles Tilly. (2015). *Contentious Politics*. Oxford: Oxford University Press.
- Thomasson, Amie. (forthcoming). "A Pragmatic Method for Conceptual Ethics." In *Conceptual Ethics and Conceptual Engineering*, ed. A. Burgess, H. Cappelen, D. Plunkett. Oxford: Oxford University Press.
- Tilly, Charles. (2001). "Mechanisms in Political Process." *Annual Review of Political Science* 4: 21-41.
- Wright, L. (1973). "Functions." *Philosophical Review* 82: 139-168.
- Yablo, Stephen. (2008). "Beyond Rigidification: The Importance of Being *Really* Actual.' In *Thoughts*, Oxford: Oxford University Press, pp. 171-190.
- _____. (2014) *Aboutness*. Princeton: Princeton University Press.
- Yalcin, Seth. (2016) "Belief as Question Sensitive." *Philosophy and Phenomenological Research* <https://doi.org/10.1111/phpr.12330>
- Zawidzki, Tadeusz. (2013) *Mindshaping: A New Framework for Understanding Human Social Cognition*. Cambridge, MA: MIT Press.