

Going On, Not in the Same Way

Sally Haslanger
shaslang@mit.edu

1. Introduction

I don't know what concepts are, or even if there are any (Machery 2009). So it feels awkward to set out to write a paper specifically on conceptual engineering. I suppose I have myself to blame, however, for I have suggested more than once that valuable projects within philosophy can be *ameliorative*, more specifically, that we should seek not only to elucidate the concepts we have, but aim to improve them in light of our legitimate purposes (Haslanger 2012, especially Ch. 6; see also Haslanger 2017). In several works, I have argued for ameliorative accounts of *gender*, *man*, *woman*, and *race*; or, in the language of concepts, for ameliorated concepts of *gender*, etc.

When I said all this (starting in the early 1990s), however, I was not thinking much about background philosophical work on concepts, and didn't have the language to express clearly what I now think is plausible. Although I was aiming to challenge traditional conceptual analysis by suggesting that our concepts might be improved by empirical or pragmatic considerations (in what I called, on one hand, *descriptive* projects, and on the other hand, *ameliorative* projects), I left it unclear whether I took the project to be semantic, metasemantic, or neither.

The fact is, I've always been much less interested in what our terms mean than in what in the world is worth talking about (Haslanger 2012, Ch. 16), though of course, these two issues are related (more on this below). It is worth noting, however, that I did not claim, at least initially (in Haslanger (2012, Ch. 6)), that my accounts of *gender* or *race*, for example, capture what we *ordinarily* mean by the relevant words, and explicitly appropriated the terms for my purposes (allowing that there are other purposes and contexts in which my accounts would be inadequate, even morally questionable).

At the time, I was interested in the idea that a critical theory should be *emancipatory*. Raymond Geuss (1981) had questioned what such a demand might involve, and how it might be achieved. Start with the idea that we all participate in social structures and enact social practices that are unjust, but most of the time this is not obvious to us, even when we are the ones disadvantaged. We get up in the morning and do our best to get our kids to school and ourselves to work on time, not thinking about the racialized school to prison pipeline, the exploitation of janitors who have cleaned our offices overnight, or our own enactment of the gendered division of labor. Our practices and our identities typically present themselves to us in ways that mask the broader system.

A critical theory challenges the understandings – understood broadly to include beliefs, concepts, scripts, narratives, norms, emotional and perceptual dispositions, and such – that motivate and appear to justify our unjust practices; it can also offer resources to think and act differently, with better epistemic and normative tools. To notice how the existing practices and structures depend on distorted understandings can itself be liberating, in a sense, for we can begin to frame new intentions, explore different forms of agency, and take on new identities. This enhances our

autonomy. But it doesn't make us free. Emancipation requires also that we change the unjust practices that structure our lives, and this requires more than thinking about oneself differently.

My own reading of Beauvoir (1989/1949), Frye (1983), Wittig (1993), MacKinnon (1987, 1989), and others, had been liberating; these works disrupted my understanding of what it is to be a woman and prompted self-reflection on my investment in gender and race categories. I vividly recall the moment when I read Wittig, "'To refuse to be a woman, however, does not mean that one has to become a man.'" (Wittig "One is Not Born a Woman," 49). The obvious truth of this blew my mind, but I had never before considered the possibility it describes. Why not? Such moments are something philosophical work can make available, and together with social activism, can contribute to social justice. Or so I believe. I saw my own theoretical project as continuous with theirs.

My work on gender was published at a particular historical moment when feminists were questioning whether we could continue to talk about women and suggested that we were entering a post-feminist era (Butler 1990, Riley 1988). I was seeking tools to resist this move. However, I see now that my accounts of *woman* and *man* problematically excluded some trans women from being counted as women and trans men from being counted as men, and failed in a number of other ways (Bettcher 2009, 2012, 2014, 2016 (and comments); Jenkins 2016; Kapusta 2016). But I still believe that the categories my accounts point to are theoretically and politically important and deserve attention, whatever we choose to call them. One lesson here is that linguistic choices that might be emancipatory at one moment, or for some individuals, or in response to a certain threat, may be inadequate in a broader context and even deepen other forms of oppression. This is not a new lesson and is to be expected. Although I gestured at these risks in early work, I didn't do enough to guard against a number of linguistic and political harms that could have been foreseen.

As time went on, I saw that amelioration could be combined with a kind of externalism that would enable me to claim that the disruptive accounts I proposed might provide a better account of what we *actually* mean by certain terms than what we take ourselves to mean (2012, Ch. 13, 14; see also Saul 2006). In attempting to capture this tension within our linguistic practices, I've suggested (2012, Ch. 2, Ch. 13) that we should, as philosophers, not get hung up trying to capture our manifest meanings (what we think we are talking about, or what we think we are doing with our words), but should aim to capture our operative meanings (what we actually do with our words, what they actually pick out). This was a way, I thought, of combining descriptive analysis of the sort proposed by post-Quinean scientific essentialists with the sort of critical, ameliorative project I was committed to. Although early scientific essentialists (Kripke 1980; Putnam 1975, et al) focused on *natural kind terms* and thought that only *natural* scientists could be relied on to find the essences of things, I aimed to broaden the reach of their projects to include not only social science, but critical social theory more broadly. Manifest meanings, as I understood them, are a potential site of ideology, for they can mask or distort how and whether our terms track kinds and sometimes contribute to the construction of them.

Interestingly, this shift in my understanding of my own accounts altered the political import of my project: emancipation, it seems, involves at least two moments. One moment is negative: we need

to understand the failures of our current practices; another moment is positive: we need to suggest better alternatives. My early accounts of *gender*, *race*, etc. might be considered useful in a negative moment to illuminate the exclusionary assumptions embedded in our use of certain social kind terms; however, they fail to offer adequate replacements (Jenkins 2016). This suggests that, more generally, efforts to provide tools for emancipation are not only context-sensitive, but may require both a disruptive moment that target a set of practices, and also what we might call a visionary moment that gives us resources to create something better.

2. Concepts?

But, one might ask, how is this narrative connected to conceptual engineering? What role do concepts play in the project? As said before, I do not have an account of concepts. But it may be helpful to identify some of my background commitments as I would articulate them now. I am anti-Fregean about meaning, content, and semantics: 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 we do represent the world – propositions are abstract entities that carry information and are, in a sense, 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 do utterances and mental states get the content they have? This is a project for *metasemantics*.

Metasemantics investigates how linguistic practices and conventions link utterances (occurring in response to others, in different parts of the world, and in other possible worlds) with propositions. It also concerns the ways in which mental states, whether linguistic or not, process and carry propositional and sub-propositional information. Some philosophers and psychologists are apt to suggest that concepts do an important part of this work. When I believe, for example, that the cat is on the mat, I have a concept of *cat* and *on* and *mat*, and various logical connectives, and I combine them in a thought. Another way to think about what’s happening is that I have a set of capacities for processing information: capacities for attention, categorization, interpretation, memory, language, inference, affect, and the like, and these capacities organize inputs in ways that represent information discursively. Propositions – that, as mentioned before, can be understood just as sets of possible worlds – are encoded in a variety of different ways, and some propositions I have easy and direct access to, whereas others I have more mediated or inferential access to.

Some of the mental capacities we have are hard-wired. However, humans and some other non-human animals have tremendous perceptual, cognitive, and affective flexibility that enable them to adapt to a variety of settings. For humans and other social animals, adaptation is deeply connected to coordination. Learned mechanisms of coordination require selective attention to public phenomena that serve as signals in response to which we do our part. A red light might take up only a tiny space in our visual field, but drivers are highly responsive to it, for failure to see it or

respond appropriately may be a matter of life or death. Language is one form of public information exchange that gives us a basis for coordination, but language itself depends on more basic capacities we have for picking up information from others, and from our environment, and sending it to others (Zawidzki 2013). In responding to and transmitting information, we develop predictable patterns of behavior that others come to expect; these patterns, when upheld by the coordination group, constitute practices. Non-human animals have such capacities well.

For example, dogs and humans coordinate. One crucial task for this trans-species coordination is the timing of the opening and closing of outside doors. Dogs need to go out to “do their business” and usually cannot open the door for themselves; they need humans to do it for them. In our house, we have a bell hung by the door. When Sparky wants to go outside, he rings the bell. We come to the door and let him out. Sparky had to be taught to use the bell. The bell does not have a “natural meaning” (Grice 1957), but it has a meaning in the ecosystem of our home. In response to his need, Sparky rings the bell, expecting that we will come open the door; we hear the door and expect Sparky to be waiting near the door to go out. The bell does not have linguistic meaning, but it has, what I call, social meaning. The ringing bell provides information that we – Sparky and the other family members – are able to access due to a process of learning from each other in an effort to coordinate. Our capacities for attention, interpretation, categorization, etc. have adjusted to take in this information and act on it in ways that have come to be expected. Not all social meanings are about coordination, but it is plausible that the capacities that make social meaning possible originated in the need for coordination. Humans, though, are able to take delight in social meanings for their own sake and use them to develop cultures that have lives of their own. (Balkin 1998; Zawidzki 2013)

We might say that if one develops a sophisticated set of capacities that enables them to process certain kinds of information, say, about X s, then they have the concept of X . Consider Yalcin:

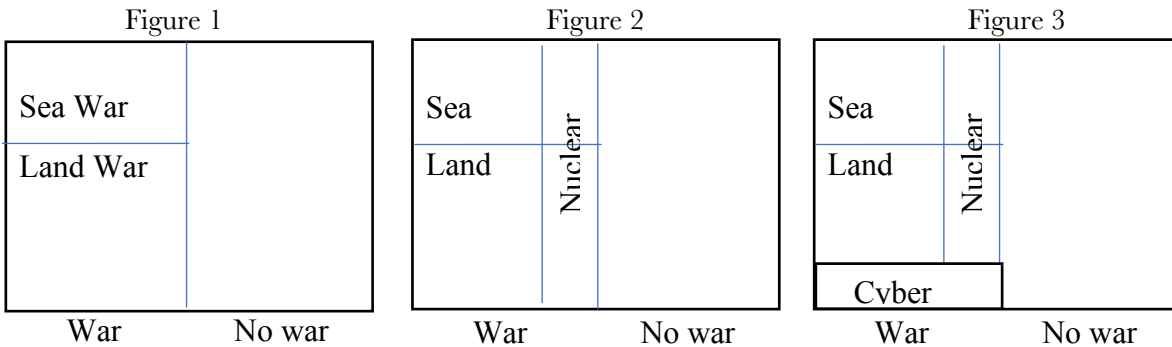
To possess a concept is to have an ability to cut logical space in a certain way, to distinguish possibilities in terms of the sorts of things that answer to the concept....A concept determines a matrix of distinctions. For example, the concept/subject matter BACHELOR corresponds to the partition of 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. 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)

The content of the concept is a partition of logical space. From a psychological point of view, however, *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 certain 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.

For example, we may have the same concept of *cat* – the informational content of the concept *cat* is the same for each of us – but our possession of it occurs in somewhat different ways so that certain

inferences are more direct for me than they are for you, or that I am more ready to apply the concept than you. Or it may be that because you know more about *cats*, you have a sensitivity to different kinds of *cat*, so your partition of logical space is more fine-grained.

Consider the concept of war. “William III of England believed that England could avoid war with France. Did he believe that England could avoid nuclear war with France?” (Yalcin 2016, 12) The content of William III’s concept of war might be represented as a division between war (land or sea), and non-war (see Figure 1). Logical space is divided into regions that contain worlds with sea wars, worlds with land wars, and worlds with no wars (considered timeless); worlds (like ours) can occur in more than one partition of space because we have had both land wars and sea wars. Utopian worlds, perhaps, have no wars. However, once we learn about the possibility of nuclear war, we can carve the space of possibilities in a more fine-grained way (see Figure 2). And as we learn more about war, and as military technology and tactics evolve, we might want to not only add complexity, but also redraw distinctions that seemed exhaustive before (See Figure 3).



William III could not draw the distinctions between kinds of wars that we draw. In fact, William III couldn’t even imagine nuclear war or cyber war. Our concept of war is more fine-grained than his and we have knowledge about more kinds of war. That is to say, we can draw distinctions between kinds of wars that he was unable to draw. But that doesn’t mean that we don’t, in an important sense, employ the same broad concept when he, and we, think and talk of war.

These tools can be used, also to clarify the way in which Aristotle and I mean the same thing by ‘water.’ The informational content of ‘water’ for both of us is the same, though he may have a rather simple way of differentiating kinds of water, e.g., potable water and non-potable water (or, given that it is Aristotle, likely he would have many ways of distinguishing kinds of water!). We differentiate, e.g., “pure” water (= H₂O), tap water, bottled water, distilled water, fluoridated water, polluted water, etc. Our concept of water is more fine-grained (and/or differently grained) than his.

There comes a point, however, when the differences between speakers and thinkers suggest that sameness of content has not been preserved. We also find terminological stretches, e.g., we now talk of the “war on drugs” and the “war on terror.” Are these really wars, or does such usage take us too far? How do we proceed when there is controversy over what war, say, really is? How do we adjudicate our disagreements? And how are these controversies related to amelioration?

In order to avoid confusion in what follows, I'll use the term 'distinction' and 'distinguish' or 'classification' and 'classify' for the linguistic/conceptual acts of noting or marking differences, and the terms 'difference,' 'differentiate,' or "division" for the ontological basis for distinctions when they divide the world, i.e., we distinguish objects that are different; our distinctions aim to capture what differentiates the objects we're interested in, we draw distinctions along divisions.

Suppose we start with this background picture: mental and linguistic representations have informational content; the *informational content* of a concept is a partition of logical space that divides possibilities.¹ To *possess* a concept is to have some cluster of capacities and mechanisms for using that grid of possibilities at some level of resolution, i.e., for making the distinction(s) in question, and processing and storing the relevant information.

On a model of this sort, one way to ameliorate would be to improve our understanding of the informational content of the concept. This would be an epistemic amelioration, but one that might shift what the speaker/thinker takes to be the relevant informational content. Individually and collectively, we use concepts without having a very solid grasp of their informational content. We may not be able to apply the concept to remote possibilities, and there may be gaps in our judgments about cases. We might say that in such cases we need to refine our concept based on a broader or deeper knowledge of the phenomenon, e.g., by gaining insight into logical space at a more fine-grained resolution (as in the war example). Or it might involve an improvement in the mechanisms by which we access the informational content, e.g., to have more reliable or illuminating access to the content by different modes of presentation: those who have been actively engaged in war have a different access to the informational content than those who only imagine or read about it.

Another way to ameliorate would be to adjust the informational content of the concept itself, i.e., to change what partition of logical space it represents. Although this, perhaps, is what the term 'conceptual amelioration' suggests, it is unclear how to spell this out. The problem seems to be this: consider two partitions of logical space D and D^* . Our concept of *bachelor*, say, has D as its content – a particular set of all and only possible bachelors. Suppose an ameliorator comes along and suggests that, instead, the concept of *bachelor* 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^* is fine and dandy to serve as the content for a concept of *bachelor*, but that would be a different concept from ours because *our concept* has D as its content. If we accept the ameliorator's new content D^* , then it would seem that we would not be improving our concept of bachelor, but adopting a new concept because concepts have their informational content essentially.

Note that both of these approaches assume that the informational content of our concept is *given*. In the first case, we are looking to improve our grasp of it; in the second case, we are considering whether to adjust it. But we still don't have a metasemantic account of what makes it the case

¹ 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). For our purposes, it isn't crucial which option we take.

that a concept has a particular content. If I think that a bachelor is must be “on the make” or available for marriage, and you don’t, what determines whether I’m right or you are? I will consider this question in the next section, however, it might be useful to have some examples before us. Water, gold, and jade have been standard examples in the literature about natural kinds. Marriage and family provide examples of controversial social kinds. Gender and race are examples where there is disagreement over whether they are natural or social.

- Is water essentially H₂O? Is H₂O the informational content of our thought and talk about water?
- According to the dominant understanding of *marriage* prior to 2004 in the United States, marriage is a legal and/or religious status restricted to one man and one woman. A marriage between two men or two women was, for many, unintelligible. Same-sex marriage is now legal in the United States, and in some religions, it is fully accepted. Both the institution of marriage and our understanding of it have changed. A consequence of this is that not only how we think of families, but how families are constituted has changed. How should we understand these changes? Have we improved our concept of *marriage* (or of *family*) or adopted a new concept?
- According to Kant, races are groups of humans who have evolved to have certain distinguishing physical and psychological traits, that “resist further transformation.” So, for example, the native American, having had to endure the extreme cold, suffers from a “half-extinguished life power” (Bernasconi et al 2000, 17). And the Negro, because he has benefited from the rich land of Africa, is “strong, fleshy, and agile. However, because he is so amply supplied by his motherland, he is also lazy, indolent, and dawdling.” (17) Whites, however, have diverged least from the original form and the “noble blond form” characterized by its “tender white skin, reddish hair, and pale blue eyes” that inhabited the northern regions of Germany, is the strongest. This form itself does not constitute a race, but only a lineage within the white race. However, “This stock would have gotten on well enough to persist as a race if the further development of this deviation had not been so frequently interrupted by interbreeding with alien stocks.” (20) Social constructionists about race reject Kant’s claims about race and propose that race is a social category (Mills 1997; Haslanger 2012, Ch. 6; Taylor 2004; Mallon 2004, 2006; Hardimon 2003, 2017; Jeffers 2013). Are social constructionists improving our concept of race or introducing a new concept? What is the relationship between Kant’s conception of race and theirs?

3. “Conceptual Analysis”²

On the model under consideration, thought and talk about water, for example, has a certain informational content. When Aristotle points to the sea and says ‘That is water,’ and when I point to the sea and say “That is water,” we are both saying something true, and what makes it true is that the sea is composed of water. What Aristotle and I represent – the content of our utterance – is the same, even if how we represent it is different. But what *are* we talking about? What is water?

² Some paragraphs in this section also appear in my chapter, “What is Race? Tracing its Socio-Political Reality,” and my “Replies,” both to appear in Glasgow et al (forthcoming).

There are a variety of ways we might go about answering questions of the form “What is X?”. In some cases the question is about a particular thing before us. For example, if, pointing to a small wiggly thing in the corner, I ask, “What is that?” I will probably want someone to help me figure out the species of insect it belongs to, as determined by entomology. But when we ask “What is X?” not about a particular (silverfish, for example), usually a better way to put the question is: *What it is to be X?*

Why might we ask such questions? Two kinds of contexts come to mind. One is when there is a conflict over whether something is, or is not, X. For example, I say that mold is a plant, and you disagree. We then look up what it is to be a plant and discover that plants are living organisms that produce their own nutrients through photosynthesis. This is *what it is to be a plant*. So I’m wrong. Molds are not plants, they are a kind of fungi. Cases of this kind don’t actually need to involve conflict. They may be more a matter of uncertainty: is a mold a plant or an animal? (Neither, we discover, because molds don’t satisfy what it is to be a plant or what it is to be an animal; figuring this out this may also require knowledge of what it is to be a mold.) Another context for asking such questions is when we want to explain why something characteristically behaves in a certain way. For example, I want to understand why the trees in my yard are dropping their leaves, and you explain what it is to be deciduous, which may, in turn, involve an explanation of what it is to photosynthesize, etc. The twofold role of *what it is to be X* – settling uncertainty over cases and seeking explanation – are both important in answering such questions.

This suggests that to answer questions about *what it is to be X*, we must situate Xs in a broader frame of understanding. The questions carry implicit contrasts, e.g., is a mold a plant (or an animal or a fungi)? Why is my tree dropping its leaves (rather than keeping them all winter)? Explaining what it is to be X, characterizes the phenomenon in ways that provide explanatory links to related phenomena.³ But it is not always clear what sort of explanation is called for, or what phenomena are relevantly related.

When considering water, or silverfish, or other natural phenomena, it would be, at the very least, odd to answer the questions by consulting our linguistic intuitions. Our judgments about when to use the term ‘silverfish’ don’t tell us what a silverfish is. However, there are a variety of “What is X?” questions that many philosophers seem to think can be answered by discovering the meaning of the term(s) substituted for X, as determined by our disposition to apply the term(s) in question, e.g., ‘knowledge,’ ‘moral worth,’ ‘justice,’ ‘a person,’ ‘causation.’ In some of these cases, one might think that this a priori methodology is warranted because the boundaries of these kinds depend in some way on us and our practices. Perhaps moral worth, justice, personhood, and the like, don’t exist independently our judgments of what counts as moral worth, justice, and personhood. So, of course, we should at least begin by investigating our judgements and putting them in order. (This is more plausible in some cases than in others, e.g., the answer would have to be more complicated in cases such as ‘causation’ or ‘intrinsic property.’)

³ I want to allow here that there are different kinds of explanations, e.g., explanations in terms of composition, function, structure, etc. See Garfinkel (1981) on explanation as a response to contrastive questions.

But the idea that (some) philosophical kinds somehow “depend on us,” is not entirely clear; nor is it clear why our a priori (linguistic) reflections should be sufficient to provide an adequate theory of them. For example, “What is a sheriff?” Even if you are a competent user of the term ‘sheriff,’ you may not be able to tell me what a sheriff is. A full answer would presumably require information about the jurisdiction of sheriffs, what their responsibilities are, how they are chosen, etc. as determined by law. We might need to consult experts in civics to get answers (and the answers will depend on what country we are in). We can’t just depend on common sense or linguistic intuitions. But surely what counts as a sheriff depends on us – there are no sheriffs outside of a humanly constructed system of government.

In the case of ‘sheriff,’ there will be a well-defined role specified by statute, and someone who knows the relevant statutes will know the answers to our questions. But there are also social phenomena that in some sense “depend on us” but are not stipulated or planned by us. Such social phenomena range from macro-scale economic depressions, globalization, urbanization, and gentrification, to more local social practices and relations, e.g., within a town, religious congregation, or family. These phenomena call for explanation, and the social sciences (broadly construed) endeavor to provide theories that enable us to understand them, usually identifying kinds of institutions, economic relations, cultural traditions, social meanings, and psychological predispositions, to do so. The kinds in question are social kinds, in the sense that they are kinds of things that exist in the social world (and so, in some sense, depend on us). But we discover these kinds through empirical enquiry just as we discover chemical kinds through empirical inquiry.

For example, accounts of gentrification often make reference to the “urban pioneer,” sometimes characterized as artists and “bohemians” who take advantage of low rents in poor neighborhoods. Once single people who share rent enter a neighborhood, businesses (such as cafés and pubs) take interest, and landlords see opportunities to raise rents, which drives out the locals. *Urban pioneers* are a functional kind that identifies a particular role in an evolving real estate market. The term ‘pioneer’ is chosen due to the perceived parallel with pioneers who “settled” the western United States, displacing the local population. If someone were to object to the term ‘pioneer’ – perhaps thinking that it carried an overly-positive connotation – this would not undermine the explanatory claims.⁴ The adequacy of explaining gentrification by reference to singles moving into an urban neighborhood does not depend on our linguistic intuitions about applying the term ‘pioneer’ to them. The choice of terminology was intended to illuminate a parallel; if the terminological choice doesn’t work, then another term could (and sometimes should) be used as a substitute.

However, insofar as philosophical kinds such as *justice* and *personhood* “depend on us,” it is not in the sense that we stipulate what they are (like *sheriff*), or in the sense that they serve in explanations of social phenomena (like *urban pioneers*). Rather, it is something along these lines: the adequacy of our theory is not to be judged simply by reference to “the facts,” but also by its responsiveness to our prior understandings. In the case of *sheriff*, you might think that there aren’t any independent facts we’re trying to accommodate. (Oversimplifying), we simply create sheriffs and then talk about

⁴ Metaphors and analogies can play an important and even ineliminable role in theorizing and can aid in explanation. My claim here is only that the choice of terminology for the functional kinds in the proposed mechanisms of gentrification (specifically the influx of singles) is not essential to the success of the model for some purposes (though it may be for others).

them. In the case of *urban pioneer*, the prior understandings of ‘pioneer’ are not crucial to the explanation provided by the theory. But in the case of *justice*, there is something we are aiming to understand that is not simply constituted by what we say, but at the same time, our conclusions cannot float completely free of the discursive tradition in which we are aiming to understand it.

How might we explain this? Note that in the philosophical cases, we are not situated as anthropologists trying to understand the social life of the “natives.” Nor are we legislators specifying new practices. We seek an understanding of practices in which we are currently engaged as participants. The practices are not fully understood, however. And they are open-ended, revisable, possibly self-defeating. In making sense of them, we are making judgments about how to better understand what we are doing, and how then to go on. This is not primarily a linguistic exercise: we aren’t just deciding how to use existing terminology, but how to collectively orient ourselves towards the world and each other. Language provides tools to achieve this. But language is a practice within practices and is itself a proper target of philosophical inquiry: meanings are not simply constituted by what we believe, yet we are situated within a tradition of linguistic practices that have already shaped our meanings and our world; so ignoring those practices would be a mistake. We are situated inquirers, and the question is how we should go on, given where we have been, where we are now, and where we are trying to go (Wittgenstein 1958, esp. §§185-243; Kripke 1982, esp. fn. 13 pp. 18-19; Lear 1986).

4. Representational traditions: ‘Water’ as an example.

Laura and François Schroeter (2015) offer an account of meaning that situates our linguistic activities within our broader social practices. They focus on the example of ‘water,’ and suggest that to determine what ‘water’ means, we should undertake an inquiry into what water *is*. “It’s important to notice that from the first-person perspective, the object-level question ‘what is x?’ is equivalent to the metalevel question ‘what is the reference (or, more generally, the semantic value) of my term “x”?’” (2015, 419)

But how do we determine what water is? We cannot assume from the start that this is a task for the chemist, for when the chemist says that water is H₂O, she may be using the term in a technical sense, in which case it would not provide an account of what the ordinary person means by ‘water.’ But neither can we just undertake reflection on linguistic usage or common sense.

Before you explicitly reflect on the question of what water is, your own assumptions about the topic are bound to be heterogeneous, incomplete, and partially contradictory – and this heterogeneity is only exacerbated when you take your whole community’s views into account. Thus justifying an answer to a ‘what is x?’ question is nothing like slotting some missing values into an implicitly grasped formula. Your goal in rational deliberation is to find some principled way of prioritizing and systematizing your own and your community’s commitments about water, so as to identify the appropriate normative standards for evaluating the truth and acceptability of beliefs about the topic. (2015, 430)

The broad idea is this: when we deliberate about what X [water, race, freewill, moral worth, gender...] is, we have to start with *something*. In the sorts of cases we are considering, we can take ourselves to be situated within a broad representational tradition concerned with X (we are not

starting from scratch and stipulating the meaning of theoretical terms). And we may assume that the tradition has a certain epistemic ambition, so we may “take our words and thoughts to represent genuinely interesting and important features of the world - not just whatever happens to satisfy our current criteria.” (2015, 436) So scientific inquiry is also relevant, since it discloses some parts of the world that are important for many of our purposes. But where do we begin?

The Schroeters (2015, 426) give a sample of inputs to deliberation in the case of water:

- *Particular instances*: there's water in this bottle, in Port Phillip Bay, Lake Michigan, etc.
- *Perceptual gestalts*: the characteristic look, taste, odor, tactile resistance and heaviness of water.
- *Physical roles*: water's rough boiling point, its transformation into steam, its role as a solvent, the fact that it expands when it freezes, etc.
- *Biological roles*: water's necessity for the survival of plants and animals; how it's ingested; the effects of water deprivation; etc.
- *Practical roles*: the roles water plays in agriculture, transport, washing, cooking, surfing, etc.
- *Symbolic roles*: water is strongly associated with cleanliness and purity, it plays an important role in many religious rituals, etc.
- *Explanatory roles*: water has a non-obvious explanatory structure, which explains many of its characteristic roles; water is composed of H₂O.
- *Epistemology*: water is easy to spot but hard to define; our beliefs about water may be mistaken or incomplete; observation of instances of water grounds induction to unobserved cases.

Our aim is to answer to the “What is X?” question. Their project is not semantic but *meta-semantic*: what would make it the case that the informational content of ‘water’ is all and only possible instances of H₂O? The inputs just considered help us narrow down the target kind so we can investigate it further. As we proceed, we may find that some of our background beliefs are false and our theoretical efforts misguided. But what do we do with these inputs? How do we balance various considerations? Schroeter and Schroeter (2015) propose that

...ideal epistemic methods for answering 'what is x?' questions hinge on rationalizing interpretation of one's representational traditions. You need to diagnose the most important representational interests at stake in a representational tradition with 'x', and you should identify the correct verdict about the nature of x as the one that makes best sense of those interests. (2015, 430)

A rationalizing interpretation, on their view, is not determined by reports of beliefs and intentions of participants in the tradition, nor is it a causal explanation of the tradition:

From the deliberative perspective of a rational epistemic agent, the interests that are relevant to adjudicating 'what is x?' questions are those that help justify or rationalize that tradition. Ideal methods for adjudicating 'what is x?' questions don't simply construe representational practices as meeting psychologically or causally fixed representational interests. Our interpretive methods construe them as meeting representational interests that help make sense of our practices - that help construe them as having a point or rationale." (2015, 435)

In the case of ‘water,’ there are at least two candidates. One set of interests served by our attitudes towards and talk of water are explanatory, another set is practical. These two interests may come apart, e.g., our practical interests do not require that we identify water with H_2O , for liquids that are mostly H_2O but contain other ingredients (harmless trace chemicals, fluoride) are fine for most purposes (drinking, bathing, swimming...). However, scientific inquiry enables us to explain the properties of water – and how it can actually serve our practical interests – by reference to its chemical structure (431). This divergence of possible interpretations of what’s at stake in the tradition leaves us with two candidate answers to “What is water?” and so the meaning of ‘water.’ Water is H_2O , or water is the watery stuff found in lakes and rivers (etc.). It might appear that this leaves the term ‘water’ is ambiguous, or perhaps with no determinate meaning.

On the Schroeters' (2015) view, however, there is a best interpretation of the representational tradition, where the scope of that tradition is determined by commitment to de jure sameness of reference and shared linguistic and epistemic practices (428). What I mean is not just a function of what I think water is, or any old interpretation of our representational tradition: I can get the meaning wrong if I don’t do justice to the interpretive task. For example, if I decide that, given our interests and collective uses of the term, water is the alcoholic beverage also known as ‘beer,’ I would be wrong. I would have failed to capture a reasonable interpretation of our representational tradition.⁵ But I could also be wrong if I miss what is worth talking about:

As rational epistemic agents, we normally take our words and thoughts to represent genuinely interesting and important features of the world - not just whatever happens to satisfy our current criteria. When asking about the nature of water (or free will, color, etc.), we don't assume that we (or our community as a whole) already implicitly know the right answer. (2015, 436)

We postulate ambiguity or opt for an error theory only as a last resort. In the case of ‘water,’ H_2O wins because it organizes and explains the relevant inputs to deliberation. So the informational content of ‘water’ (and similar words in other languages), is a partition in logical space on which chemical composition is the differentiating feature.

5. Conceptual Amelioration

This gives us a helpful model for understanding at least some conceptual amelioration. We are part of a complicated representational tradition with various threads and various purposes for employing the concept X . Although we are adept at using the concept in many contexts, its exact informational content is obscure. Perhaps we find ourselves in controversy, or in a situation where more precision or understanding is needed. We engage in reflection on the representational tradition and find that certain ways of going on – ways of interpreting our past practices projecting them into the future – requires adjustment in our judgments about what counts as an X , or what is true of X s.

⁵ The Schroeters seem to hold that there is one best interpretation of the tradition. I am not convinced. More on this in what follows. See also Botchkina and Hodges 2016.

Water is perhaps not the best example when we are thinking about amelioration. In the case of ‘water,’ or other natural kinds, we typically begin with a set of uncontroversial cases, broadly shared interests, and little variation in usage. This makes it plausible that Aristotle and I are talking about the same thing when he speaks of ὕδωρ and I speak of water, and that we share this meaning with others in our linguistic community. Aristotle might have paused when he first saw snow to consider whether it is also ὕδωρ. But taking snow and liquid water to be of the same kind was easily justified by further observation, e.g, melting snow becomes water. The suggestion that Aristotle and I are in the same linguistic community is grounded in the similarity of our practices with respect to watery stuff and the naturalness of the kind. The fact that he couldn’t conceive of hydrogen or oxygen, and that we live in different periods in history with different practices and technologies, means that we have different ways of partitioning the sub-kinds of water (as described in §2), but this does not undermine the claim that the informational content of ‘water’ and ‘ὕδωρ’ are the same.

But consider the concept of *family*. In certain periods in the United States, the bionormative nuclear family has been (and in some current contexts is) definitive of what families are. This might suggest that in those social milieu’s the concept of family is something like this:

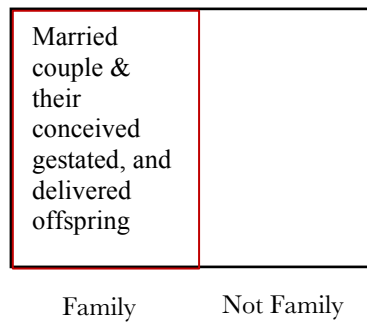



Figure 4

However, when we reflect on the representational tradition within which we think and talk of families in the United States, we find that the practical, symbolic, explanatory, and legal definitions of family don’t correspond to this division in logical space. Changes in reproductive technology, discoveries of different ways of organizing domestic life in different cultures and times, and pressures from social movements, make it plausible that the line between family and not-family lies at a different place.

Looking at the phenomena broadly from our current point of view, we might say that to have a *family*-concept is (roughly) to have a cluster of mechanisms for processing information about the coordination of domestic life. Married couples and their “fully” biological offspring (MCBOs) may constitute one kind of family, but not *what it is to be a family*. A better way of understanding the representational tradition would include such MCBO families only as a sub-group of all families, e.g., we should find ways to include families made by adoption, donated gametes, families with single parents, same-sex parents, unmarried parents, and extended families of various kinds (see

Figure 5 – this is not intended as a complete or even adequate representation of the logical space of families!).

Figure 5



Suppose, however, that Albert has a concept of family that seems to correspond to the partition of logical space in Figure 4. Albert has *a* concept of family. Albert may resist calling same-sex couples raising children a “family.” But unless your processing of information is such as to enable you to coordinate in our milieu around *Xs*, it would be plausible to say that you don’t have *our* concept of *X*. “Our” concept of family is embedded in our practices and our laws, our forms of intimacy and love.

We may aim to ameliorate Albert’s concept of family by pointing out to him that “our” representational tradition supports a broader understanding of what a family is. This would seem to offer him a better grasp of what he means, i.e., an epistemic amelioration. Note that he may resist this intervention. He may have a different interpretation of our representational tradition, or he may not care what the rest of us think and instead choose to consider a different representational tradition as his own, e.g., a particular religious tradition. Nevertheless, if he aims to think and speak about families along with those in the broader community of the United States, then amelioration of his understanding of what families are, i.e., updating how he understands the informational content of ‘family,’ is called for.

However, this may not seem adequate to capture what some want in the idea of conceptual amelioration. Yes, some people are out of sync with the broader representational tradition of which they are a part, and if they want to conform to that tradition, they should adjust their understandings of what they are talking about. But the idea of conceptual amelioration is especially valuable when we are cutting against the grain – there are times when the community seems to be (or has been) committed to a particular understanding that should be resisted. Conceptual amelioration should just be a matter of demanding that the outliers conform to the dominant understanding of what we are talking about. This calls for a more robustly normative sense of amelioration.

Return to the concept of *family*. On one hand, there is a descriptive project of characterizing all the possible ways of organizing domestic life. But on the other hand, there is a normative project. We know that some ways of organizing family life are morally problematic and other forms lead to various kinds of social dysfunction. One normative axis of evaluation is functional: Because we process certain kinds of information mainly for the purpose of coordination, we can judge the

adequacy of our ways of doing so by reference to how well the coordination works. Better and worse concepts of family might then be evaluated in terms of how well they achieve what domestic life is for. Another axis of evaluation may be broadly moral. Perhaps Albert believes that there is something morally problematic about non-MCBO families, so we should organize domestic life by creating and supporting MCBOs; all other forms are somehow defective and so should be discouraged. On his view, only MCBOs are *real* families. He is not denying that we call many domestic arrangements ‘families.’ He would agree that the issue is how we should go on, but thinks that we are currently going in the wrong direction. He prefers a more narrow way of understanding what a family “really” is.

To understand the tension between the descriptive and the normative aspects of amelioration, it is useful to consider what Joshua Knobe and Sandeep Prasada call “dual character concepts” (2011, also Knobe et al. 2013; Leslie 2015). Note that some kind terms allow for a distinction between being a good exemplar of the kind and being a true exemplar. Usually these two evaluations go together. A *good* scientist is a true scientist; a true musician is a good musician. However, these evaluations seem to be based on connected but distinct criteria. They argue that there is a set of concepts – the dual character concepts – that

...are represented via both (a) a set of concrete features and (b) a set of abstract values that the concrete features are seen as realizing. These two representations are intrinsically related, but they are nonetheless distinct, and they can sometimes yield opposing verdicts about whether a particular object counts as a category member or not. (2011, 2965)

One of their paradigm examples concerns the concept of *scientist*. On the one hand, we might characterize a scientist in terms of the sorts of things they do, their qualifications, their job description, etc.; on the other hand, we might characterize a scientist (roughly) in terms of their intellectual virtues. They note that these two characterizations of a scientist might come apart (2011, 2965). One might say of a person who works in a lab but is dogmatic and does the minimum required each day:

(1) There is a sense in which she is clearly a scientist, but ultimately, if you think about what it really means to be a scientist, you would have to say that she is not a scientist at all.

Or one might say of a person not employed in a lab or academic institution and who “has never been trained in formal experimental methods but who approaches everything in life by systematically revising her beliefs in light of empirical evidence” (2011, 2965):


(2) There is a sense in which she is clearly not a scientist, but ultimately, if you think about what it really means to be a scientist, you would have to say that she truly is a scientist.

Other examples that scored high as dual character concepts include: “Friend, Criminal, Love, Mentor, Comedian, Minister, Theory, Boyfriend, Artist, Argument, Teacher, Poem, Soldier, Sculpture, Art Museum, Musician, Mother, Rock Music, Scientist, Novel” (Knobe et al, 2013, 256).

What is going on in these cases? Knobe et al argue that there are two kinds of normativity playing a role in the case of dual character concepts. On one hand, we can evaluate the extent to which one meets certain conditions for being a member of the kind in question, e.g., a *good* musician plays

fluently, has advanced skills (even though they may have little creative spark). On the other hand, we can evaluate the extent to which one exemplifies certain abstract values, e.g., a *true* musician lives to make music, and does so creatively and with passion (even though they may not have fabulous skills). The goodness v. the trueness of a musician are different dimensions of evaluation.

Knobe et al suggest that the duality is illuminated by a comparison between dual character concepts and natural kind concepts. Natural kind concepts seem to be committed to there being “hidden causes” that give rise to or cause the distinctive characteristics of the kind, e.g., zebra DNA gives rise to characteristic zebra stripes; in the case of dual character concepts, abstract values (being a devoted seeker of empirical knowledge) are characteristically realized by a pattern of features or behaviors (working in a lab) (2013, 244).



However, in the case of natural kinds, when the “hidden cause” and the superficial features come apart, e.g., if something is made to look like a zebra, but doesn’t have zebra DNA, judgments of membership in the kind follow the “hidden cause.” In the case of dual character concepts, however,

...people appeared to employ two distinct sets of criteria. When a given object met one set of criteria but not the other, participants tended to say that it was a category member in one sense but was not a category member in another sense. As such, the experiment provided evidence that dual character concepts provide two bases for categorization. (2013, 248)

Nonetheless, they hypothesize that there is an important parallel between natural kind concepts and dual character concepts. In both cases, the unity of the ordinary conditions for being a member of the category is explained by reference to the supposition of a “hidden” factor, such as an inner essence or an abstract value. For example, what unifies the conditions for being a *scientist* and a *good scientist* depends on the abstract values purportedly embodied by the kind. They suggest further:

We have seen that some concepts are unified through hidden causes (natural kind concepts) and others through abstract values (dual character concepts), but perhaps these are just two of the many possibilities, and there are also yet other kinds of concepts that are unified in quite different ways. For example, there might be concepts in which all of the concrete features are unified in that they all tend to make an object suitable for the same basic function (e.g., the concept COMPUTER). People might then associate these concepts with both (a) a list of concrete features and (b) the more abstract notion of the relevant function. (If so, such concepts would be like the dual character concepts studied here in that they would provide two bases for categorization, but they would be unlike dual character concepts in that they would not provide two bases for normative judgment.) (2013, 255).

A plausible hypothesis is that we draw distinctions for different purposes.⁶ Sometimes we categorize for the purpose of explanation, sometimes to capture what's of value, sometimes to identify a functional role (and as Knobe et al suggest, possibly for other purposes as well). There can be controversy, then, over not only what conditions must be met in order to be included in a category, but also the point or purpose of the category, and what gives the category its unity. My suspicion is that inquiry into categories sometimes rightly privilege the (paradigm) instances and allow our understanding of the purpose to adjust; and sometimes we are rightly invested in the purpose and reconsider the instances. Note, however, that the partition of logical space that actually serves the purposes of our distinction (once we figure out the purpose), may be very different from what it was before, based on our judgments of membership, or even based on our conception of what's at stake. (Baseball doesn't cease to be baseball when we change the rules, or "improve" the ball to allow more home runs so that it attracts more fans.) This provides a basis for allowing ameliorative accounts to be, in some sense, a way of improving and not just replacing our concepts.

Drawing on these considerations, it might seem that controversies over the concept of family involve several different issues. Is the family a "natural kind"? Is it a functional kind? Is it a dual character concept to be understood in terms of a background abstract value? Different accounts of *family* will see it as serving different purposes and guided by different desiderata.

Let's return to Albert's concept of family (figure 4). It would seem likely that he would grant that, under current circumstances, same-sex families count as families "in some sense," e.g., legally, but aren't *true* families. That is because they don't realize certain abstract (family-) values. On this interpretation, Albert has a dual character concept of family. This would allow him to say that the functionalist model (figure 5), which demarcates families in terms of their role in coordinating domestic life, captures "a sense" of family, but only a subset count as "true" families. To capture the disagreement between Albert and the functionalist, then, we could see it as concerned with the proper function of families, and the values that should be realized if that function is fulfilled. This, I assume, is not something that can be established simply by looking at our past practices, or past judgements of memberships, for our past practices may all be terrible. Normative discussion is needed in order to decide how to go on.

6. Conclusion

There is much more that needs to be discussed at this point. However, I've argued that there are two ways to think about conceptual amelioration on a model according to which the content of a concept is to be understood as informational content, i.e., a partition of logical space. On the one hand, we can ameliorate our *understanding* of the relevant space of worlds – both what worlds it includes and how they should be sub-divided. Such improvement in our understanding may be based on advances in empirical knowledge and technological advances (think of war, water, and offspring created by artificial reproductive technologies), on a reinterpretation of the history of the

⁶ I don't mean to suggest that we have and develop concepts for explicit and intended purposes; we don't. We are socialized into the local conceptual/linguistic scheme and no one designed it. I do think, however, that concepts play a functional role in systems of communication and coordination, and I have this sort of purposiveness in mind. See also Haslanger (forthcoming).

concept, or on information gained by new perspectives that have access to the content from different modes of presentation.

On the other hand, we can ameliorate not simply by reinterpreting our past practice but by *correcting* it. It is not plausible to me that the “best” interpretation of our representational tradition concerning race is an interpretation according to which race is a social category. Nor am I convinced that *marriage* has always included in its informational content worlds in which there were same-sex marriages. I’m not even sure how those judgments would be adjudicated. But on my view, this doesn’t matter. Marriage, as I understand it, is a functional and normative category: it includes those domestic partnerships that promote a well-functioning and just society; it includes same-sex marriages. I don’t make this claim based on what we have meant all along, but on what we should have meant. And what we should mean going forward, at least for now. And I hope that those who follow us will not be held back by our current efforts to make sense to each other.

Some of our concepts are organized around values. Others are organized around functions. Some are organized around both. This is because we have an interest in carving logical space in order to coordinate with each other, to draw distinctions that serve our purposes as social beings and to realize our values. The best way to do this changes as we develop new technologies and as we come to appreciate new and different values. When social change happens, there is likely to be controversy and disagreement about how to extend the concepts we’ve been using to do the work we now need them to do. Such changes should be acknowledged as such, and should not be held hostage to what we have thought we were doing all along, and how to continue that. Our conceptual frameworks should be forward-looking and give us the tools to envision and create better lives together.

Acknowledgements (incomplete):

Thanks to Ekaterina Botchkina, Alexis Burgess, Alex Byrne, Herman Capellen, Esa Diaz-Leon, Jerome Hodges, Karen Jones, Ari Koslow, Mari Mikkola, Deborah Mühlebach, David Plunkett, Greg Restall, Mark Richard, Laura Schroeter, Francois Schroeter, Stephen Yablo, Seth Yalcin....

Works Cited

- Balkin, J.M. (1998) *Cultural Software: A Theory of Ideology*. New Haven: Yale University Press.
- de Beauvoir, Simone. (1989/1949) *The Second Sex*. Trans. H.M. Parshley. New York: Vintage.
- Bernasconi, Robert and Tommy Lott. (2000) *The Idea of Race*. Indianapolis: Hackett Publishing.
- Bettcher, Talia. (2009) “Trans identities and first-person authority”, in *You've changed: Sex reassignment and personal identity*, Laurie Shrage (ed.), Oxford: Oxford University Press, 98–120.
- _____. (2012) “Trans women and the meaning of ‘woman’.” in *Philosophy of sex: Contemporary readings (sixth edition)*, Nicholas Power, Raja Halwani, Alan Soble eds.), New York: Rowan & Littlefield, 233–250.
- _____. (2014) "Feminist Perspectives on Trans Issues", *The Stanford Encyclopedia of Philosophy* (Spring Edition), Edward N. Zalta (ed.), URL = [<https://plato.stanford.edu/archives/spr2014/entries/feminism-trans/>](https://plato.stanford.edu/archives/spr2014/entries/feminism-trans/).
- _____. (2016) *Precis of and comments on “Katharine Jenkins' "Amelioration and Inclusion: Gender Identity and the Concept of Woman," Pea Soup* January 27, 2016:

- <http://peasoup.typepad.com/peasoup/2016/01/ethics-discussions-at-pea-soup-katharine-jenkins-amelioration-and-inclusion-gender-identity-and-the-.html>
- Botchkina, Ekaterina and Jerome Hodges. 2016. "Objectivity and Conceptual Change."
Unpublished manuscript.
- Butler, Judith. (1990) *Gender Trouble*. New York: Routledge.
- Frye, Marilyn. (1983) *The Politics of Reality*. Freedom, CA: Crossing Press.
- Garfinkel, Alan. (1981) *Forms of Explanation*. New Haven: Yale University Press.
- Geuss, Raymond. (1981) *The Idea of a Critical Theory*. Cambridge: Cambridge University Press.
- Glasgow, Joshua, Sally Haslanger, Chike Jeffers, and Quayshawn Spencer. (Forthcoming) *Four Views on Race*. Oxford: Oxford University Press.
- Grice, H. Paul. 1957. "Meaning." *The Philosophical Review* 66(3): 377-388.
- Hardimon, Michael. (2003) "The Ordinary Concept of Race." *Journal of Philosophy* 100(9): 437-455.
- _____. (2017) *Rethinking Race: The Case for Deflationary Realism*. Cambridge, MA: Harvard University Press.
- Haslanger, Sally. (2012) *Resisting Reality: Social Construction and Social Critique*. Oxford: Oxford University Press.
- _____. (forthcoming) "Cognition as a Social Skill." *Australian Philosophical Review*.
- Jeffers, Chike. (2013) "The Cultural Theory of Race: Yet Another Look at Du Bois's 'The Conservation of Races,'" *Ethics* 123 (April 2013): 403-426.
- Jenkins, Katharine. 2016. "Amelioration and Inclusion: Gender Identity and the Concept of Woman." *Ethics* 126(2): 394-421.
- Knobe, Joshua & Sandeep Prasada. (2011) "Dual Character Concepts", In *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*. Boston, MA: Cognitive Science Society.
- Knobe, Joshua & Sandeep Prasada, and G.E. Newman. (2013) "Dual Character Concepts and the Normative Dimension of Conceptual Representation." *Cognition* 127(2): 242-257.
- Kripke, Saul. (1980) *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- _____. (1982) *Wittgenstein on Rules and Private Language*. Cambridge MA: Harvard University Press.
- Lakoff, George. (1999) "Cognitive Models and Prototype Theory." In Eric Margolis and Stephen Laurence, eds. *Concepts: Core Readings*. Cambridge, MA: MIT Press, pp. 391-423.
- Lear, Jonathan. 1986. "Transcendental Anthropology." In Philip Pettit and John McDowell (ed.) *Subject, Thought and Context*. Oxford: Oxford University Press, pp. 267-298.
- Leslie, Sarah-Jane. (2015) "'Hillary Clinton is the only man in the Obama administration': Dual Character Concepts, Generics, and Gender" *Analytic Philosophy* 56 (2):111-141.
- Machery, Edouard. (2009) *Doing Without Concepts*. Oxford: Oxford University Press.
- MacKinnon, Catharine. (1987) *Feminism Unmodified*. Cambridge, MA: Harvard University Press.
- _____. (1989) *Towards a Feminist Theory of the State*. Cambridge, MA: Harvard University Press.
- Mallon, Ron. (2004) "Passing, Traveling and Reality: Social Constructionism and the Metaphysics of Race," *Noûs*, 38 (4): 644-673.
- _____. (2006) "Race: Normative, Not Metaphysical or Semantic," *Ethics*, 116 (3): 525-551.
- Mills, Charles. (1997) *The Racial Contract*, Ithaca, NY: Cornell University Press.
- Pérez Carballo, Alejandro. (2016) "Structuring Logical Space." *Philosophy and Phenomenological Research* 92(2): 460-491.
- Putnam, Hilary. (1975) "The Meaning of 'Meaning'". In H. Putnam, *Mind, Language and Reality*. Vol 2 of *Philosophical Papers*. Cambridge: Cambridge University Press.

- Riley, Denise. (1988) *'Am I That Name?': Feminism and the Category of 'Women' in History*. Basingstoke: Palgrave-McMillan.
- Saul, Jennifer. (2006) "Philosophical Analysis and Social Kinds II." *Proceedings of the Aristotelian Society Supplementary Volumes* 80: 119-143.
- 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. (1998) What Might Nonconceptual Content Be? *Philosophical Issues* 9: 339– 352.
- Taylor, Paul. (2004) *Race: A Philosophical Introduction*, Cambridge, UK: Polity Press.
- Wittgenstein, Ludwig. 1958. *Philosophical Investigations*. Oxford: Basil Blackwell.
- Wittig, Monique. (1993) *The Straight Mind*. New York: Basic Books.
- Yablo, Stephen. (2014) *Aboutness*. Princeton: Princeton University Press.
- Yalcin, Seth. (2016) "Belief as Question Sensitive." *Philosophy and Phenomenological Research*
- Zawidzki, Tadeusz. (2013) *Mindshaping: A New Framework for Understanding Human Social Cognition*. Cambridge, MA: MIT Press.