

# Conceptual Engineering and Semantic Deference

Joey Pollock

ConceptLab, Department of Philosophy, University of Oslo

---

Many ameliorative projects aim at moral goods such as social equality. For example, the amelioration of the concept MARRIAGE forms part of efforts to achieve equal rights for the LGBT+ community. What does implementation of such an ameliorated concept consist in? In this paper, I argue that, for some ameliorated concepts, successful implementation requires that individuals eschew semantic deference, at least with respect to relevant dimensions of the concept. My argument appeals to consideration of the aims of conceptual engineers engaged in this type of ameliorative project: they seek conceptual change in order to contribute to the dismantling of oppressive social structures, institutions, and systems of belief. I argue that, for such aims to be achieved, it must be the case that individuals who come to endorse the concept do so for the right reasons—because they have gained an understanding of why the ameliorated concept is morally preferable to its ancestor. Once they have acquired such reasons, however, they are no longer correctly described as semantically deferential; they will treat moral reasons to employ the concept as overriding of semantic considerations.

*Keywords:* conceptual engineering, semantic deference, amelioration, implementation

---

## 1. Introduction

Many of the concepts we possess are products of the social world we live in; they can also shape the ways in which we represent and make sense of this world. As such, it is not surprising that concepts can reflect, perpetuate, or obfuscate those features of our social world that are ethically problematic. Concepts may more or less explicitly embed assumptions that contribute to oppression; or they may simply fail to be politically useful ways to represent a particular phenomenon or group. Critical scrutiny of the concepts we inherit from our language community, then, is of great importance in efforts to achieve social equality. Where a concept is found to be problematic,

*Corresponding author's address:* Joey Pollock, Department of Philosophy, University of Oslo, Postboks 1020 Blindern, 0315 Oslo, Norway. Email: joeykpollock@googlemail.com.

it should be transformed into something better, or perhaps eliminated from our repertoire altogether. This is the job of the conceptual engineer: conceptual engineers are interested in what concepts we *should* have, where those that we do in fact employ may fail to serve our moral, political or theoretical purposes.

My interest in this paper is with conceptual engineering projects that aim at social equality. For example, engineering of the concept `MARRIAGE` forms part of efforts to achieve equality for the LGBT+ community. The engineering of `MARRIAGE` is an instance of ameliorative analysis: the concept is improved or replaced with a better one. Once the initial ameliorative analysis has been carried out, the resultant concept must then be implemented in a target community. My aim in this paper is to argue for one condition on the successful implementation of some kinds of ameliorated concept. I will argue that (for some concepts) successful implementation requires that the individual who acquires the ameliorated concept eschews semantic deference with respect to this concept (or certain dimensions of this concept). The reason for this is that, in these cases, successful implementation requires that individuals treat moral reasons to use concepts as overriding of semantic considerations.

The paper proceeds as follows. In Section 2, I introduce ameliorative analysis. In Section 3, I introduce some assumptions and preliminary remarks concerning concepts and semantic deference. In Section 4, I argue for the relationship between successful implementation of ameliorated concepts and semantic deference.

## 2. Engineering concepts for social equality

An ameliorative analysis begins by identifying a defect in some existing concept. A proposal is then developed for improving the concept to remove the defect.<sup>1</sup> The result is a new concept or, perhaps, an improved version of the old one.<sup>2</sup> In what follows, I will refer to the original concept as the “deficient” or “ancestor” concept, and the concept that results from amelioration as the “ameliorated” or “new” concept. I will focus on the amelioration of the concept `MARRIAGE` as my central example. The ameliorative analysis of `MARRIAGE` has both a conceptual and a legal dimension. On the conceptual level, the project is to improve the concept such that, where it previously could apply only to partnerships between a man and a woman, it should

<sup>1</sup> Simion (2018) argues that concepts do not need to be defective in order for us to want to improve them through revisions.

<sup>2</sup> Not all authors would describe ameliorative analysis this way. Haslanger (2006), for example, suggests that ameliorative analysis can be a process of revealing the nature of existing concepts.

instead be broadened such that it is inclusive of partnerships between individuals of any sex/gender, such as same-sex partnerships. One of the reasons to seek this improvement is that it can aid in the effort to secure marriage equality; that is, to secure marriage as a legal right for those in the LGBT+ community to whom it is currently denied. However, one could also be in favour of the conceptual broadening of marriage even if one thought that the *legal* institution of marriage should be abolished altogether (Card 1996, Young 1997). The amelioration of MARRIAGE aims at various goods intended to contribute to equality for the LGBT+ community. For example, the legal institution of marriage, in many countries, comes with significant economic and legal benefits that are not otherwise available to partners in a relationship.<sup>3</sup> Aside from these benefits, though, the ameliorative project might contribute to combating further dimensions of heterosexism. Callahan (2009), for example, points out that the benefits that are available to married partners have been symbolic of the fact that the state holds only heterosexual relationships and families as legitimate and worthy of protection; and Calhoun (2002) argues that the denial of marriage equality contributes to a cultural conception of lesbians and gay men as defective citizens. I will return to Calhoun's argument in Section 4.

As demonstrated in this sort of example, an ameliorative project can aim at a range of effects in different communities, where these effects are hoped to contribute to achieving and sustaining social equality for oppressed groups. Few, if any, of these effects will be achieved by merely ameliorating a concept, however. Rather, the success of an ameliorative project requires that the concept be *implemented* in a target community or communities. Ameliorative projects alone are surely not sufficient to achieve social equality for any given group. However, they can contribute to this end; and it is only when a concept is implemented in a target community that it can start to have some kind of impact on things like social structures, institutions, the belief systems of individuals, and so forth. Successful implementation within a target community might require uptake and retention of the concept amongst some or all of its individual members, depending on the aims of the ameliorative project.<sup>4</sup> In this paper, I will restrict my focus to the question of what successful implementation demands of an individual who acquires the ameliorated concept. I argue that, for some concepts, successful implementation (with respect to an individual) requires that the individual who acquires the concept does not semantically defer with respect to its application (or relevant dimensions of its application). Before arguing for this

<sup>3</sup> Card (1996) argues that these benefits render marriage an unjust and dangerous institution that should be abolished.

<sup>4</sup> Thank you to Alex Davies for pressing me to be clearer about this issue.

claim, I will first say a bit about how I am understanding both concepts and semantic deference.

### 3. Concepts and semantic deference

There are a great many views as to what concepts are and how they are individuated.<sup>5</sup> However, theories of concepts, despite their differences, tend to agree (more or less) on a range of desiderata that specify the sorts of roles that concepts play in an individual's cognitive life and interactions with others. Concepts are the components of thought contents and, as such, are supposed to be central to understanding phenomena such as categorisation, learning, reasoning, communication, and intentional action. Whether one thinks that concepts are the physical realizers of conceptual roles in an individual, or abstract objects determined by patterns of use across an entire community, one will typically take one's theory to be answerable to such desiderata. Beyond this rough characterisation of the theoretical role of concepts, I wish to remain largely neutral regarding which theory of concepts is correct.<sup>6</sup> However, I will make a few assumptions that rule out some approaches. Firstly, I assume that, whatever concepts are, they are the sorts of things that can (perhaps only with great effort) be deliberately altered in predictable ways. Not all approaches allow for this. Cappelen's (2018) Austerity Framework, for example, claims that the sort of deliberate engineering of concepts described by those engaged in ameliorative analysis is not possible.<sup>7</sup> For the purposes of this paper, I am going to set aside Cappelen's worries. I restrict my focus to theories that claim that conceptual engineering is a process that is under our control and within our ken. Secondly, I assume that the amelioration of an existing concept is a matter of an individual, or group of individuals, developing a *new* concept—one that is not type-identical to the old.<sup>8</sup> I think this assumption will not be too controversial. It is plausible that conceptual engineers can simultaneously possess both the ameliorated concept and its deficient ancestor. Many conceptual engineers will be experts on both concepts—understanding how the ancestor concept is in fact

<sup>5</sup> For a range of approaches to content individuation, see (Rapaport 2002), (Farkas 2008), (Chalmers 2002), (Burge 1979), (Jackman 2005), (Millikan 1989), (Fodor 1987).

<sup>6</sup> For the reader who finds this attempt at neutrality unhelpful, I think my argument can be run using talk of inferential roles.

<sup>7</sup> Cappelen states his framework in terms of the extensions and intensions of expressions rather than concepts, but his points about control and inscrutability can be carried over to mental content.

<sup>8</sup> In claiming that this is a new concept, I do not mean to suggest that it amounts to a change of topic. For accounts of topic continuity, see (Cappelen 2018), (Prinzling 2018), (Sawyer 2018).

employed in a community, whilst nonetheless believing that it *should not* be employed.<sup>9</sup>

Aside from those views ruled out by my assumptions, the argument that I wish to make is compatible with many views of concepts. However, the terminology best used to frame it might differ slightly depending on whether one takes concepts to be prototypes or inferential roles, or internally or externally individuated, or what have you.<sup>10</sup> That individuals are disposed to defer with respect to many of the concepts and expressions they employ is a datum that any theory of concepts should seek to be consistent with; and we can theorise about when individuals ought to defer without first settling the question of what concepts are. I will have more to say about semantic deference in Section 4. For now, let us say that semantic deference is the phenomenon whereby an individual accepts correction with respect to her application of an expression or concept from those she deems to be more expert language users. The relationships of deference between individuals form a structure that links lay speakers with (those who are perceived to be) the most competent, or “expert”, speakers at the top. As Burge writes:

We may imagine a vast, ragged network of interdependence, established by patterns of deference which lead back to people who would elicit the assent of others. [...] To put it crudely, a person counts as among the *most competent* if he or she would be *persuasive* to other competent speakers in the use and explication of the language. (Burge 1986, 702)

Semantic deference is a phenomenon that is most commonly discussed in connection with social externalism. This is the view that concepts are individuated, in part, by an individual’s social environment (Burge 1979, Goldberg 2007).<sup>11</sup> For a social externalist, an individual’s dispositions to defer allow her to share in an inventory of communal concepts (Burge 1979, 114). However, other approaches to concepts can allow that there is such a phenomenon as semantic deference. Content internalists, for example, will agree that subjects often accept correction from those they deem more competent; what they deny is that these dispositions enable subjects to possess communal concepts that they do not understand. For the internalist, typically,

<sup>9</sup> Thank you to Ishani Maitra for this point.

<sup>10</sup> I will note when such differences of terminology are required as they become relevant.

<sup>11</sup> Burge’s considered view often seems closer to a form of physical externalism; he claims that the physical environment plays a more decisive role in content individuation (1986; 1988; 1979, fn 2). Sawyer (2018) endorses a pure social externalism for semantic content, but not for thought content. She claims that semantic content is determined by what the community would agree to after reaching reflective equilibrium in debate. Concepts, on the other hand, are individuated by relations to real entities and kinds in the environment.

having deferential dispositions indicates that subjects desire to acquire those concepts that most closely approximate the concepts of the experts. In what follows, I argue that, for some concepts, successful implementation of these concepts (for an individual) requires that the individual who acquires them is no longer semantically deferential with respect to their application.

#### 4. Implementation

A full account of successful implementation is beyond the scope of this paper. However, before getting to my argument, it will be useful to identify some further conditions on implementation that are less controversial.

##### 4.1 Understanding and endorsement

One such condition (depending on one's theory of acquisition) is that an individual must not only acquire but also understand the ameliorated concept, at least along certain dimensions.<sup>12</sup> With respect to the sort of example under consideration, the social changes at which the ameliorative project aims require that individuals form new beliefs (and jettison old beliefs) about the nature of the phenomenon represented by the relevant concept, which track the proposed conceptual change. For example, agents need to believe that it is possible for marriages to be formed between partners of any sex/gender. Forming these beliefs will often require a significant degree of conceptual understanding of the ameliorated concept, as well as understanding of some of the concepts with which the central concept interlocks. It may sometimes be the case that full understanding, or conceptual mastery, is required for implementation. However, there may also be cases in which an individual need only understand the ameliorated concept along certain relevant dimensions.<sup>13</sup> For ease of exposition, I will often talk of an individual's need to understand "the concept", but I do not intend to claim anything stronger than that she must understand this concept along dimensions relevant to the ameliorative project.

Another condition that might be required (depending on one's account of conceptual understanding) is that the individual "endorse" the ameliorated concept. On some views, even a condition as strong as concept mastery would require only that an individual fully understand a concept; it would not require that she actually apply it to the group or phenomenon for which

<sup>12</sup> This way of framing the point is more congenial to externalist views. Many internalists will claim that one cannot possess concepts that one does not understand; as such, acquisition and mastery do not come apart.

<sup>13</sup> Thank you to Jeff Engelhardt for this point.

it was designed.<sup>14</sup> Returning to our thought from earlier: implementation requires that individuals form and maintain certain kinds of belief about the objects and phenomena represented by the ameliorated concept. Mastering a concept requires, at most, that a subject grasp (perhaps implicitly) the conditions under which the concept purports to apply. As noted above, the position that many conceptual engineers may occupy is that of mastering *both* the ameliorated concept and its deficient ancestor, whilst believing only one of these ought to apply. Those resistant to the proposed amelioration may possess both concepts whilst endorsing only the ancestor. If an individual acquires the ameliorated concept but does not use it in place of its ancestor; or, worse, if she masters the new concept only so that she may more effectively design barriers to its uptake in a community, I think it is clear that the implementation has not been successful. Successful implementation, then, should require that the individual endorse the concept over its ancestor, not merely that she understand it.<sup>15</sup>

To recap: for the implementation of an ameliorated concept in an individual's conceptual repertoire, so far, we have it that she must understand the concept (at least along certain relevant dimensions) and she must endorse the concept over its ancestor. In what follows, I will argue that this alone would not be sufficient for successful implementation for some kinds of ameliorative project. In some cases, the individual must not only endorse the new concept, she must do so for specific reasons: she should endorse the concept, not because others do, but because she thinks it is the right thing to do, morally speaking. I will first explain why the moral reasons for employing a concept should be considered important to implementation in the kinds of ameliorative project under consideration; I will then argue that the condition requires giving up on semantic deference (with respect to the concept, or relevant dimensions of it).

## 4.2 Implementation and deference

The kind of reasons that are relevant to my argument are motivating reasons: reasons that explain why an individual endorses the concepts that she does.<sup>16</sup>

<sup>14</sup> Some views of concept mastery incorporate endorsement into their conditions on mastery. See (Rabin 2018) for an argument against this kind of approach. Rabin argues for a "recognition" view of mastery according to which mastering a concept is a matter of correctly recognising certain core elements as governing the use of the concept (without necessarily endorsing these elements).

<sup>15</sup> Ameliorative projects may not always aim to replace a deficient concept with just *one* new concept. There may be cases where a plurality of ameliorated concepts is useful (See Brigandt and Rosario, forthcoming).

<sup>16</sup> I do not mean to suggest that there are no such things as normative reasons for endorsing a concept. Conceptual engineers who aim at moral goods can be wrong about which con-

For many concepts, the reason that we use them is just that other people do. Conceptual engineers, as theorists engaged in the project of critiquing the concepts we inherit from others, will typically have different sorts of reasons for using ameliorated concepts—ameliorated concepts are used because they are thought to better serve various purposes than their inherited ancestors. Certainly, the reasons for ameliorating a concept are important to theorists and activists; but why should we think that it is also important that such reasons are taken on board by the individuals who adopt the new concept? Roughly, the idea is this: as noted earlier, in ameliorative projects of the sort under consideration, engineers are motivated by the pursuit of moral goods—they are interested in contributing to the dismantling of oppressive social structures, institutions and systems of belief, and replacing them with those that will promote and sustain social equality. When, in implementing an ameliorated concept, engineers seek to transmit their own motivating reasons to the individuals who adopt the new concept, it is because acceptance of these reasons is partly constitutive of the sorts of changes that they wish to effect—individuals who adopt the new concept ought (if the project is to be considered successful) to be appropriately motivated by specific moral considerations. Calhoun (2002) is explicit about the importance of transmitting such reasons. She writes that, “The moral significance of extending rights is to a large extent a function of the sorts of arguments that get culturally circulated in the process of extending rights” (Calhoun 2002, 109).<sup>17</sup> Calhoun is here focused on the issue of extending legal rights, rather than on ameliorating concepts. I do not wish to attribute to her any particular view regarding conceptual engineering and implementation. However, I do think that her remarks regarding arguments for extending rights are equally applicable to the implementation of ameliorated concepts.<sup>18</sup>

To see why the circulation of reasons is important to implementation, it will be useful to look at Calhoun’s argument for the pursuit of marriage equality in more detail.<sup>19</sup> Calhoun considers a number of potential moral arguments in favour of marriage equality and its priority on the political

cepts are morally preferable and, thus, depending on one’s view of normative reasons, an individual may have a moral reason not to employ these concepts in this second sense. For the distinction between normative and motivating reasons, see, e.g., (Dancy 2002), (Smith 1994).

<sup>17</sup> The same point is made by Richardson-Self (2015, 127). Richardson-Self evaluates a range of arguments in favour of marriage equality, including Calhoun’s.

<sup>18</sup> One might also think that altering concepts is required in order to achieve uptake of particular justifications for extending rights.

<sup>19</sup> Calhoun focuses on the importance of marriage equality for lesbians and gay men. I will follow her own presentation of her argument, but I think similar considerations are relevant to other kinds of non-heterosexual orientations.



agenda, many of which she rejects. One option that she considers is that marriage equality might help to undermine sexism and male dominance (Calhoun 2002, 115ff)—for example, by promoting the idea that gendered husband/wife roles are inessential to marriage and thus encouraging heterosexuals to adopt more egalitarian relationship models. Calhoun thinks that this is not the right argument to circulate, for several reasons. One of those reasons is that, in framing the injustice of same-sex marriage bars in terms of gender inequality, this approach fails to acknowledge that heterosexual domination is an axis of oppression in its own right: she argues that the forms of oppression suffered by the LGBT+ community are distinct from gender-based oppression and cannot be explained solely in terms of it. Calhoun's preferred argument runs as follows. She argues that being deemed fit for marriage is culturally bound up with what it means to be a true citizen. This is because the dominant cultural conception of marriage treats it as a pre-political institution: it is not something that is created by the state; instead the very possibility of civil society depends on people entering marriages and forming families.<sup>20</sup> The state protects, promotes and regulates marriage because it is thought to play this foundational role. In this ideology, people who are deemed fit to marry and form families have a special status—they are necessary for the foundations of society. In contrast, those who are deemed unfit for this role have a lesser status. Bars on same-sex marriage thus reflect and perpetuate the ideology that lesbians and gay men are inessential or lesser citizens who are unfit for family life; Calhoun describes them as *displaced* to outside civil society. We should pursue marriage equality, according to Calhoun, because bars to same-sex marriage play a central role in the oppression of lesbians and gay men by displacing them in this way.<sup>21</sup>

In relation to ameliorative analysis, one way we can describe these alternative arguments is as underpinning different concepts (or clusters of concepts), each consistent with the conceptual change described at a more superficial level.<sup>22</sup> An ameliorative strategy based on Calhoun's work would require that an individual endorse, not just any inclusive marriage concept, but the particular concept-cluster that she is proposing—for it is the endorse-

<sup>20</sup> Calhoun here is focusing on conceptions of marriage in the USA, as expressed in the 1996 Defence of Marriage Act. She is not suggesting that this is a conception that we ought to endorse.

<sup>21</sup> Calhoun (2002, 128ff) notes that, beyond achieving marriage equality, there are different revisions to the cultural conception of marriage that might be adopted. For example, we might also reject the conception of marriage as a pre-political institution (although she thinks that there are some reasons to retain this conception).

<sup>22</sup> I do not mean to suggest that such arguments are reducible without remainder to the endorsement of concept-clusters.

ment of this concept-cluster that is needed to contribute to undermining heterosexism in the way that she envisages. I think that this is not all that should be required by the ameliorative project, however. For a Calhoun-inspired ameliorative project to succeed, the reason that an individual ought to endorse this concept-cluster is, not only because she believes that the cultural conception of lesbians and gay men as defective citizens is central to their oppression, but also that she believes that combatting such oppression is the morally right course of action and is motivated to choose such a course of action. One could perhaps convince an individual to employ an inclusive marriage concept (or concept-cluster) for other reasons. For example, perhaps you can convince an individual to do so simply by paying her a lot of money; but, in doing so, you will not have done very much to address her heterosexist attitudes. This is not to say that there could be absolutely no benefit to getting individuals to employ ameliorated concepts for alternative reasons;<sup>23</sup> but doing so would fall far short of the aims of an ameliorative project that sought to more comprehensively address the individual's morally problematic system of beliefs and dispositions.

This first claim was a claim about the importance of the moral reasons offered for endorsing the ameliorated concept (or concept-cluster): given the aims of the ameliorative project, individuals should endorse the ameliorated concept for the engineer's moral reasons because doing so is required to overwrite specific existing problematic beliefs and dispositions with those that will contribute to promoting and sustaining social equality for the relevant group. I will now argue that, for the success of the sort of ameliorative project under consideration, individuals ought no longer be semantically deferential with respect to the concept they endorse (or relevant dimensions of it). The claim here is not that the subject need not defer because she has mastered the concept. As Burge explains, even subjects who master a concept (and thus, in fact, need no correction) are not thereby rendered non-deferential towards, or non-reliant on, the norms of their community (Burge 1979, 84). Even where an individual masters a concept, it can be true of her that, had expert use been different, she *would* accept correction. What I will argue is that an individual ought to treat her moral motivations for employing the new concept as overriding of semantic considerations, and this renders her no longer semantically deferential with respect to that concept.

Consider an individual who has adopted the ameliorated concept,  $MARRIAGE_A$ , for the sorts of reasons presented by Calhoun. As already explained, the aim of the relevant ameliorative project is to contribute to overwriting

<sup>23</sup> Indeed, there may be cases in which there are significant benefits that can be gained from convincing individuals to endorse ameliorated concepts even when they do so for alternative reasons. Thank you to Alex Davies and Jeff Engelhardt for this point.

specific beliefs and dispositions in the individual who adopts the concept. Given this kind of aim, if the implementation of the ameliorated concept is to be considered successful then, in coming to endorse *MARRIAGE<sub>A</sub>*, it should not be the case that the individual has signed up to mean *just whatever the experts mean* by “marriage”, as she might in acquiring a concept like *ARTHRITIS*. We should not consider the implementation to be successful if the individual is willing to relinquish her problematic beliefs, but only to the extent that doing so does not require her to use concepts or expressions in ways that conflict with the practices of her colinguals. A subject who is willing to let semantic considerations outweigh moral considerations in choosing what concepts to employ does not seem to be appropriately motivated to do the right thing. She should be using the new concept because she has reason to think that it is the morally right (or morally better) concept to employ, not because she has flagged a semantic expert as someone to whom she will defer with respect to future usage. If merely corrected in the manner that a semantic expert corrects a lay speaker, she ought to resist this correction. Successful implementation, then, requires, not just that individuals have moral reasons to employ a concept, but that they treat their moral reasons for employing a concept as overriding of competing considerations.

One thing that an objector may claim at this point is that, whilst implementing an ameliorated concept might require inculcating some degree of resilience in an agent, this resilience is compatible with the claim that the agent is nonetheless semantically deferential. We might describe the situation thus: there are subjects who, with respect to some concepts, are disposed to change their use with little persuasion, and there are subjects who (for various reasons), are disposed to take significantly more persuasion. But where a person sits on this spectrum, with respect to a given concept, is not relevant to the question of whether they should be described as semantically deferential: a subject can be highly resistant to changing their practice, requiring much evidence and persuasion, and yet still qualify as semantically deferential with respect to the relevant concept.<sup>24</sup> Indeed, one of Burge’s examples involves this kind of subject: one who believes sofas are religious artefacts and is highly resistant (though not immune) to semantic correction (Burge 1986, 263ff).

How should we respond to this objection? Given that individuals *can* eschew communal standards, the question the objector must answer is this: when a subject is deferential with respect to a concept, what disposition(s) is it that she possesses (with respect to that concept) that she lacks when it comes to concepts with respect to which she is non-deferential? In what follows, I will explain what I think is Burge’s answer to this question and

<sup>24</sup> Thank you to Sarah Sawyer for helpful discussion of this objection.

argue that it supports my claim that subjects will be no longer semantically deferential when they treat moral reasons to employ concepts as overriding.

Burge describes the differences between semantically deferential and non-deferential subjects in the following sorts of ways:

The subject's willingness to submit his statement and belief to the arbitration of an authority suggests a willingness to have his words taken in the normal way—regardless of mistaken associations with the word. Typically, the subject will regard recourse to a dictionary, and to the rest of us, as at once a check on his usage and his belief. When the verdict goes against him, he will not usually plead that we have simply misunderstood his views. (Burge 1979, 101)

A person born and bred in the parent community might simply decide (unilaterally) to follow the usage of the regional dialect or even to fashion his own usage with regard to particular words, self-consciously opting out of the parent community's conventions in these particulars. In such a case, members of the parent community would not, and should not, attribute mental contents to him on the basis of homophonic construal of his words. (Burge 1979, 114)

to a fair degree, mentalistic attribution rests not on the subject's having mastered the contents of the attribution, and not on his having behavioral dispositions peculiarly relevant to those contents, but on his having a certain responsibility to communal conventions governing, and conceptions associated with, symbols that he is disposed to use. (Burge 1979, 115)

What the above passages suggest is that, at least as Burge is conceiving of it, the semantically deferring subject is one who takes a particular kind of reason as relevant to her use of language and concepts: she is responsive to the ways others use terms in her community and takes herself to be answerable to this community's standards. The evidence that she will treat as relevant to her usage is of a particular kind: it is the kind found in a dictionary or encyclopaedia, or that can be gained from expert speakers.<sup>25</sup> On this picture, a subject can indeed be resistant to correction. But her reasons for resistance will reflect this kind of evidence: she may be skeptical that a dictionary is up-to-date, or that the alleged experts are semantically competent. When the semantically deferential subject is resistant, then, it is because she is skeptical of claims as to the common standard. The non-deferential subject, in

<sup>25</sup> Some versions of social externalism (including perhaps Burge's) allow that experts can be wrong in their explications of a concept. I think this approach can be accommodated in the present discussion. On such a view, the ultimate check on usage that the semantically deferential subject will be responsive to is the nature of the phenomenon or kind represented by the expert's concept. A non-deferential subject, in contrast, would not take this evidence as relevant to her use of concepts.

contrast, rejects the common standard (with respect to some word-form). She is thus not amenable to the same type of persuasion as the deferential subject: convincing her that communal norms favour a particular usage of some word-form is not relevant to her decision as to how to employ or understand the concept that she expresses with that word-form.

Let us return to our example of the resilient subject who refuses to revise her MARRIAGE<sub>A</sub> concept in response to semantic correction. Is she resilient due to skepticism as to the semantic competence of her corrector? That is, does convincing her require dictionaries, encyclopaedias, or expert speakers? The answer is no. There are two reasons for this. The first is that, given that she has mastered the relevant dimensions of the ameliorated concept, anyone suggesting that she ought to change her use along one of these dimensions must be proposing that she switch to employing a different concept.<sup>26</sup> Perhaps competing engineers have developed a concept (or concept-cluster) that they think is superior; or perhaps the original engineers have decided to revert back to the ancestor concept. The resilient subject will not accept semantic correction under these circumstances. In fact, the question of *which* concept to endorse (of several competing options) does not seem like a question that semantic experts (*qua* semantic experts) are in a position to answer: they will tell you how to understand the concepts that you do in fact possess, but they are not really in the business of telling you which concept to endorse in the first place.<sup>27</sup> The most they may claim in favour of *endorsing* their preferred concept is that most language users (or perhaps the most competent language users) endorse this concept. This sort of reason will not override the resilient subject's moral reasons for endorsing the original ameliorated concept.

The objector may well agree with the above but claim that it fails to address the manner in which the resilient subject is semantically deferential. This brings us to the second point. As mentioned above, Burge characterizes semantic deference, in part, in terms of counterfactual dependency: even in cases of actual mastery, if the subject's social environment had been different, she would accept correction. Thus, although she needs no correction in the actual world, if she *would* accept correction counterfactually, this is reason to think that, in the actual world, she is semantically deferential. Would the resilient subject accept semantic correction in the counterfactual sce-

<sup>26</sup> This leaves open the possibility that the subject is still deferential with respect to dimensions of the concept that are not relevant to the ameliorative project.

<sup>27</sup> This way of framing things is more congenial to externalist views of concepts. An internalist might rather say that an expert can inform you as to the properties of concepts that you want to employ, but is not really in the business of telling you which concepts to want in the first place.

nario? That is, would she treat factors such as dictionaries, communal usage, or semantic experts as relevant to her application of the ameliorated concept in the counterfactual world? Again, I think the answer is no. Whilst, in this counterfactual scenario, she might learn that an expression is not used as she thought it was in her community, she will not treat this as a reason to change her practice. In this scenario, she will not dispute what the communal standards *are*, or suggest that the experts' own understanding of the concept they employ is incorrect. Rather she will dispute what these standards *ought* to be: she has moral reasons to think that they ought to be different. The subject will even resist correction from the conceptual engineers from whom she learnt the concept in such a counterfactual scenario. Suppose, for example, that the subject came to endorse the ameliorated MARRIAGE<sub>A</sub> through misunderstanding an engineer's argument in the counterfactual world (where this counterfactual misunderstanding matches the Calhoun-inspired argument in the actual world). In such a case, a subject who treats moral reasons as overriding of semantic considerations will not accept mere semantic correction even from this conceptual engineer. The sort of evidence that will convince her to change her practice is, not semantic evidence, but argument regarding which concept, morally speaking, ought to be employed.<sup>28</sup> She is, thus, precisely someone who rejects the communal standards as they stand and so is not bound by them. As such, she lacks the dispositions characteristic of semantic deference.

I will note two things about this understanding of the resilient subject in closing. Firstly, whilst she will not be moved by semantic considerations, she should not be stubborn in the face of good moral reasons to change her concepts. In the present context, there are certain groups that she plausibly ought to be especially open to learning from; for example, standpoint theorists suggest that members of oppressed groups might be in a better position to gain, through critical reflection, a more accurate perspective on oppressive social relations (Hartsock 1983, Collins 2000); as such, they may be better situated to engage in fruitful engineering projects. So, the resilient subject can and should be open to changing her mind; but, if she does so, it will be because she has been given (or has herself generated) moral reasons that she finds compelling, not because of facts about how terms are used in her community or by particular experts within it. Secondly, my approach leaves open the possibility that agents can choose who to semantically defer to for moral reasons. For example, perhaps an agent, while deliberating over which ameliorated concept to employ, will flag a particular person as a moral expert on the relevant topic from whom she will accept semantic correction

<sup>28</sup> This sort of dispute might be explicitly about concepts, or it might take the form of metalinguistic negotiation (Plunkett and Sundell 2013).

on how to use certain concepts. This, I think, would amount to her accepting moral testimony regarding *which* concepts to employ and then semantically deferring to these moral experts with respect to her understanding and application of those concepts.<sup>29</sup> My approach does not suggest that an agent cannot or should not do this. Indeed, there may be good moral reasons for an individual to accept moral testimony regarding such matters. Wiland (2017), for example, argues in favour of trusting the moral testimony of those who suffer epistemic harms due to identity prejudice in their audiences. A variety of such harms have been identified in the literature (see, e.g., Collins 2000, Fricker 2007, Dotson 2011). Wiland argues that harms suffered by individuals who are silenced and/or unjustly treated as lacking credibility can be especially great when it is moral testimony concerning their own subordinated social experience that is disregarded by their audience. Thus, perhaps an individual should be especially receptive to this kind of moral testimony. I think that choosing which concepts to employ based on moral testimony can be an important stage in the moral improvement of an individual's conceptual repertoire. However, for an ameliorative project of the sort under consideration to succeed, the agent who endorses the new concept should do more than rely on moral testimony. The ameliorative project should not be considered successful until the individual achieves her own understanding of the engineer's (non testimony-based) moral reasons for endorsing the concept.<sup>30</sup>

## 5. Conclusion

In this paper, I have argued that the implementation of ameliorated concepts, in some cases, requires that individuals eschew semantic deference, at least with respect to certain dimensions of those concepts. The reason for thinking this stems from consideration of the aims of conceptual engineers engaged in ameliorative projects that aim at moral goods such as social equality: they seek conceptual change in order to contribute to the dismantling of oppressive social structures, institutions and systems of belief. The view I have proposed is that, in order for engineers to achieve their aims, an individual who adopts the ameliorated concept must not only understand and endorse this concept but must do so for the moral reasons provided

<sup>29</sup> Some authors argue that there is something puzzling or suboptimal about accepting moral testimony. For diagnoses of this puzzle, see, e.g., (McGrath 2009), (Driver 2006), (Hills 2010).

<sup>30</sup> This is not quite the same as claiming that the individual must acquire moral understanding in Hills's (2010) sense. Hills's notion of moral understanding is factive; whereas, as noted above, a conceptual engineer might be trying to implement a concept that she falsely believes to be morally superior.

by the conceptual engineer. Once they have acquired such reasons, however, they are no longer correctly described as semantically deferential; the reasons to which they are ultimately responsive will be moral rather than semantic.

### Acknowledgements

I am grateful to Alex Davies and Jeff Engelhardt for their helpful and extensive comments. Thanks also to Ishani Maitra, Sarah Sawyer, and Jon Robson for comments on earlier drafts and in conversation, and to audiences at the University of Oslo, University of Oxford, and University of Tokyo.

### Bibliography

- Burge, T. (1979). Individualism and the mental, *Midwest Studies in Philosophy* 4: 73–121.
- Burge, T. (1986). Intellectual norms and foundations of mind, *Journal of Philosophy* 83: 697–720.
- Burge, T. (1988). Individualism and self-knowledge, *Journal of Philosophy* 85: 649–63.
- Calhoun, C. (2002). *Feminism, the Family, and the Politics of the Closet: Lesbian and Gay Displacement*, Oxford University Press, Oxford.
- Callahan, J. (2009). Same-sex marriage: Why it matters — at least for now, *Hypatia* 24: 70–80.
- Cappelen, H. (2018). *Fixing Language*, Oxford University Press, Oxford.
- Card, C. (1996). Against marriage and motherhood, *Hypatia* 11: 1–23.
- Chalmers, D. (2002). The components of content, in D. Chalmers (ed.), *The Philosophy of Mind: Classic and Contemporary Readings*, Oxford University Press, Oxford, pp. 608–633.
- Collins, P. H. (2000). *Black Feminist Thought: Knowledge, Consciousness, and the Politics of Empowerment.*, Routledge, New York.
- Dancy, J. (2002). *Practical Reality*, Oxford University Press, Oxford.
- Dotson, K. (2011). Tracking epistemic violence, tracking practices of silencing, *Hypatia* 26: 236–257.
- Driver, J. (2006). Autonomy and the asymmetry problem for moral expertise, *Philosophical Studies* 128: 619–44.



- Farkas, K. (2008). *The Subject's Point of View*, Oxford University Press, Oxford.
- Fodor, J. (1987). *Psychosemantics*, MIT Press, Cambridge, MA.
- Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*, Oxford University Press, Oxford.
- Goldberg, S. (2007). *Anti-Individualism: Mind and Language, Knowledge and Justification*, Cambridge University Press, Cambridge.
- Hartsock, N. (1983). The feminist standpoint: Developing the ground for a specifically feminist historical materialism, in S. Harding and M. Hintikka (eds), *Discovering Reality: Feminist Perspectives on Epistemology, Metaphysics, Methodology and the Philosophy of Science*, D. Reidel, Dordrecht, pp. 283–310.
- Haslanger, S. (2006). What good are our intuitions? Philosophical analysis and social kinds, *Aristotelian Society Supplementary* **80**: 89–118.
- Hills, A. (2010). *The Beloved Self: Morality and the Challenge from Egoism*, Oxford University Press, Oxford.
- Jackman, H. (2005). Temporal externalism, deference, and our ordinary linguistic practice, *Pacific Philosophical Quarterly* **86**: 365–380.
- McGrath, S. (2009). The puzzle of pure moral deference, *Philosophical Perspectives* **23**: 321–44.
- Millikan, R. (1989). Biosemantics, *Journal of Philosophy* **86**: 281–97.
- Plunkett, D. and Sundell, T. (2013). Disagreement and the semantics of normative and evaluative terms, *Philosophers' Imprint* **13**: 1–37.
- Prinzinger, M. (2018). The revisionist's rubric: Conceptual engineering and the discontinuity objection, *Inquiry* **61**: 854–880.
- Rabin, G. O. (2018). *Toward a Theory of Concept Mastery: The Recognition View*, Erkenntnis.
- Rapaport, W. (2002). Holism, conceptual-role semantics, and syntactic semantics, *Minds and Machines* **12**: 3–59.
- Richardson-Self, L. (2015). *Justifying Same-sex Marriage: A Philosophical Investigation*, Rowman & Littlefield, London.
- Sawyer, S. (2018). The importance of concepts, *Proceedings of the Aristotelian Society* **118**: 127–147.
- Simion, M. (2018). The 'should' in conceptual engineering, *Inquiry* **61**: 914–928.
- Smith, M. (1994). *The Moral Problem*, Blackwell, Oxford.

- Wiland, E. (2017). Moral testimony: Going on the offensive, in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics 12*, Oxford University Press, Oxford, pp. 51–75.
- Young, I. M. (1997). Reflections on families in the age of Murphy Brown: On justice, gender, and sexuality, *Intersecting Voices: Dilemmas of Gender, Political Philosophy, and Policy*, Princeton University Press, Princeton, pp. 95–113.