

Moral Constraints on Gender Concepts¹

N. G. Laskowski

Draft, 22 March 2019

Forthcoming in *Ethical Theory and Moral Practice*, pending presentation at BSET 2019

1 Introduction

Is ‘woman’ a *sex* term that speakers use to talk about human biological anatomy?² Is it a *gender* term that speakers use to talk about a non-biological topic, such as the roles that individuals occupy in society? Mainstream use of the term suggests an affirmative answer to the former question³, while usage among trans communities and trans-inclusive feminist activists suggests an affirmative answer to the latter.⁴ Call these patterns of use the *usage data constraint*. Despite these appearances, there is some degree of pressure not to answer both questions affirmatively, given the preference among linguists to prefer the simplest account of words on which learning and communicating with them is easiest.⁵ Call this the *communicative constraint*. The guiding question of this paper is whether it is possible to provide an account of ‘woman’ that can satisfy both the usage data and communicative constraints.

¹ Thanks to Robin Dembroff, Esa Diaz-Leon, Nathan Robert Howard, Sofia Huerter, Daniel James, Jennifer Saul, and Daniel Wodak for constructive feedback on various parts of my project. Thanks also to audiences at the Workshop on The Nature and Significance of Social Kinds at the Institute for Advanced Studies in the Humanities in Essen, at the Workshop on Human Kinds at KU Leuven, at the 2018 Rocky Mountain Ethics Congress, and at the Philosophical Anthropology and Ethics Research Group seminar organized by Neil Roughley at Duisburg-Essen Universität.

² My focus is on the term ‘woman’, but everything I say applies *mutatis mutandis* to the term ‘man’.

³ “If one asks an ordinary speaker what it is to be a woman, she/he will almost certainly answer in terms of biological traits, most likely genitalia or chromosomes. The same is true of dictionaries.” (Saul 2012: 197)

⁴ “...many speakers will sometimes use “woman” as though it is not a sex term after all—this is what they do when they refer to trans women who have not undergone reassignment surgeries as “women,” a usage that seems perfectly acceptable to many of us.” (Saul 2012: 199–200)

⁵ See Dowell (2016). Saul also seems on board: “...feminists need to communicate successfully both with each other and with those who are not (yet) feminists, feminists should want to avoid large-scale misunderstandings wherever possible.” (Saul 2012: 197)

Recently, many philosophers, including Saul (2012), Diaz-Leon (2016), and Barnes (2019: 17), express varying degrees of support for a view of ‘woman’ that looks up to the task.⁶ The basic idea is that ‘woman’ has a core profile across contexts of use, thereby satisfying the communicative constraint, but the core is also sensitive to conversational standards that shift from one context to the next, thereby satisfying the usage data constraint. In Section 2, I trace the development of this *contextualist* view about ‘woman’ from Saul, who offers the view for examination, to Diaz-Leon, who defends a version of the view outright. In Section 3, I argue that the modifications Diaz-Leon makes to get around Saul’s worries with the view sap it of the flexibility that motivates contextualism about ‘woman’ in the first place.

My criticisms point toward an alternative view. In Section 4, I float the idea that ‘woman’ is ambiguous in a special way. In addition to capturing the apparent shiftiness of the term, I argue that such an ambiguity view makes better sense of the coherence and significance of some forms of trans-inclusive feminist activism. A constructive lesson that emerges from this paper is that we should not expect the semantic facts associated with terms like ‘woman’ to mirror facts about how we ought to treat each other morally.

2 Contextualism

The term ‘woman’ appears to live a double life. It seems that speakers use it as a *sex* term to talk about biological anatomy, such as in medical offices or rape crisis centers. But it also seems that speakers use ‘woman’ as a *gender* term to talk about features of individuals that involve biological anatomy less obviously, e.g. social roles or self-identifications.⁷ One high-profile reaction to this phenomenon comes from Saul (2012: 201), who puts forward

⁶ Saul offers but ultimately rejects the view, Barnes is “sympathetic”, and Diaz-Leon is all-in.

⁷ Some philosophers, including Jenkins (2016) and Wodak and Dembroff (2018), have been popularizing the idea that we need to disambiguate *gender class* from *gender identity*. Because Saul and Diaz-Leon do not make this distinction, and because the dust is far from settled on the appropriate application of it, I follow Saul and Diaz-Leon in not employing this distinction. More substantively, the core issue of this paper is whether we use ‘woman’ to discuss two recognizably separate topics concerning humans. It’s not quite clear to me that much hangs on the labels for those topics.

the following view on which ‘woman’ is sensitive to whichever conversational standards are salient in context.

Saul’s Contextualism “X is a *woman* is true in a context C iff X is human and relevantly similar (according to the standards at work in C) to most of those possessing all of the biological markers of female sex.”

This view purports to capture the conditions under which propositions expressed by uses of sentences containing ‘woman’ are true in context. I’ll pass over the first, “human” condition on the right-hand side of the statement without comment. The second condition has two parts, the first of which concerning “standards at work”, I’ll discuss when I turn to Diaz-Leon’s view in a moment. As for the second part of the second condition involving “possessing all of the biological markers of female sex,” Saul is using ‘female’ as a sex term to talk exclusively about XX chromosomes, vaginal genitalia, and the like.⁸ (Saul 2012: 215) In addition to purporting to capture the truth-conditions of ‘woman’, the view on offer from Saul looks like it has the resources to explain the term’s double life. Consider the following case.

Case 1 Lee, who has XX chromosomes, vaginal genitalia, and the like, utters ‘I’m a woman’ in a discussion concerning screenings for vaginal cancer at a medical office.

Assuming it is plausible to maintain that one of the “standards at work” is the standard of having female anatomy, the view on offer from Saul predicts that ‘woman’ contributes a constituent concerning female anatomy to the proposition expressed by the utterance of the sentence ‘I’m a woman’. That would appear to explain why ‘woman’ is being used as a sex term in Case 1. Now consider another case.

⁸ I follow Saul in both employing this terminology and not wading into the controversy surrounding it. For discussion, see Mikkola (2017).

Case 2 Harper, who has XY chromosomes, penial genitalia, and the like, utters ‘I’m a woman’ in a discussion among trans-inclusive feminist activists.

Unlike in the previous case, none of the “standards at work” plausibly involve female anatomy in Case 2. So, the view on offer from Saul predicts that ‘woman’ doesn’t contribute a propositional constituent involving female anatomy. That appears to explain why Harper uses ‘woman’ as a gender term in Case 2.⁹ The view on offer from Saul looks like it can explain the double life of ‘woman’ – it can satisfy the usage data constraint.

Nevertheless, Saul stops short of endorsing contextualism about ‘woman’ out of concern for the predictions that the view yields in the full range of cases involving trans individuals. (Saul 2012: 203) Consider the following case.

Case 3 Ash, who possess XY chromosomes, penile genitalia, and the like, utters ‘I’m a woman’ in a conversation in which the dominant number of participants are *traditionalists* who reject any non-sex involving use of ‘woman’ explicitly.

Since the “standards at work” appear female anatomy-involving, Saul would worry that Ash isn’t using ‘woman’ as a gender term – that Ash’s utterance of ‘I’m a woman’ is false. Saul would suggest that the problem with such a result is that it “doesn’t do justice to trans women’s claims”.¹⁰ (Saul 2012: 210) Saul is well aware that not everyone would share this concern. According to Saul, it’s a concern that raises difficult issues concerning whether political, moral, or more broadly normative considerations have a role to play in philosophical semantics. For the sake of argument, I follow Saul in assuming that an

⁹ Results like these lead Saul to frame the view as an improvement over Haslanger’s (2000).

¹⁰ Arguably, a version of this case in which it is one of the traditionalist’s using ‘woman’ instead of Ash (as in e.g. a traditionalist uttering ‘You, Ash, aren’t a woman’) might be one in which it is even clearer that the term is being used as sex term, since on some views, speaker intentions play a big role in determining how our words are used. Feel free to substitute such a version of the case if that seems more compelling. I am using Case 3 as it is because I don’t find the alternative kind of case to pose a clearer challenge, and I wish to maintain structural continuity with the other cases in this paper.

adequate characterization of ‘woman’ should predict that Ash is using ‘woman’ as a gender term in Case 3.

Diaz-Leon (2016) is a fellow traveler, offering a full-throated defense of a contextualist approach to characterizing ‘woman’. In particular, Diaz-Leon argues that such an approach has no trouble with cases like Case 3. In part because the dominant number of conversational participants are traditionalists in Case 3, it looks like the “standards at work” in it are sex-based, involving only XX chromosomes, vaginal genitalia, and the like. This seems to mean, as we’ve seen, that Ash is using ‘woman’ as a sex term on Saul’s contextualism. However, according to Diaz, Saul’s contextualism is stuck with this result in virtue of the *attitudes* of traditionalists. On Saul’s contextualism, Ash is using ‘woman’ as a sex term because the view includes the attitudes of traditionalists among those features determining the “standards at work” in context. In light of this observation, Diaz-Leon suggests a path forward: replace this feature with a new one. Diaz-Leon’s idea is that *moral* or more broadly *normative* considerations play a role in determining the “standards at work” in context.

Diaz-Leon’s Contextualism X is a *woman* is true in a context C iff X is human and relevantly similar (according to the standards at work in C) to most of those possessing all of the biological markers of female sex, where the standards at work are determined by the overall balance of normative considerations.” (Diaz-Leon 2016: 251)¹¹

This view looks like it has the resources to make the right predictions. There seem to be non-biological features in some of the cases that *matter*, morally speaking. For example, it is plausible to think that denying an individual’s self-identification causes some degree

¹¹ Diaz-Leon offers several statements of the view that differ subtly from one another. My presentation of Diaz-Leon’s view strives to illustrate the evolution of the view from Saul, while remaining faithful to Diaz-Leon’s suggested changes.

of distress. Such denials might be thought to exhibit disrespect or risk inculcating feelings of alienation, perhaps not unlike the way in which it might be disrespectful or alienating to insist to a theist that they aren't really a believer.¹² If so, then the balance of normative considerations might determine that a non-biological, gender-involving standard is “at work” in Case 3, not a biological, sex-involving standard. Plausibly, too, Ash is relevantly similar in their self-identification to most of those possessing all of the biological markers of female sex. Thus, Diaz-Leon's view predicts that Ash is using ‘woman’ as a gender term in Case 3.

Diaz-Leon's insight is that cases like Ash's don't reveal a problem for the kind of contextualist view that Saul proposes, *per se*. Rather, such cases bring out that the problem lies in those features determining the “standards at work” in context, i.e. the *theory of salience* baked into Saul's contextualism. Before evaluating Diaz-Leon's contextualism more closely, I am first going to head off a tempting objection to Diaz-Leon's account of Case 3 that will clarify an important feature of the view.

Again, Diaz-Leon's innovation is the claim that normative considerations play a role in determining the standards that are salient in context. This is what allows Diaz-Leon's view to predict that because self-identification matters morally, Ash is using ‘woman’ as a gender term. On important, broadly Humean traditions in moral philosophy, however, that which matters morally depends on our attitudes, like our desires. On such views, it could be that the desires of traditionalists, including their (say) desire to maintain mainstream linguistic practice, play a role in determining that which matters in Case 3. If so, then having female sex characteristics could matter most in that context. On this line of thought, Diaz-Leon's view predicts that Ash uses ‘woman’ as a gender term only on the assumption that what matters *doesn't* depend on attitudes. But Diaz-Leon is in no position to make such an assumption. Or so one might object.

Grant that what matters morally depends in some way on our desire-like attitudes. On this supposition, having female-sex characteristics matters, because traditionalists desire to

¹² See Kapusta (2016) and Wodak and Dembroff (2018).

maintain the mainstream linguistic practice of using ‘woman’ to talk about female sex characteristics. On this supposition, too, self-identification also matters, because Ash (say) desires not to be misgendered.

However, it doesn’t follow that self-identification matters *more* than having female sex characteristics. Nor does it follow that having female sex characteristics matters *more* than self-identification. For it to follow that self-identification matters more than having female sex characteristics, the desire to avoid being misgendered has to carry *more moral weight*. For it to follow that having female sex characteristics matters more than self-identification, the desire to maintain mainstream linguistic practice has to carry *more moral weight*. The claim that what matters morally depends on our desires entails neither of these further claims on its own. Thus, the question of whether morality depends on our desires is orthogonal to the question of which standard matters most in Case 3.¹³ Diaz-Leon need not worry about the tempting objection.

Nevertheless, the tempting objection does make clear that another assumption is at play in this debate. Namely, Diaz-Leon must assume that some features, e.g. self-identification, carry *some* moral weight, at least enough to outweigh whatever weight there might be in having female-sex characteristics some of the time. After all, that’s how Diaz-Leon’s view predicts that Ash uses ‘woman’ as a gender term in Case 3. In the next section, however, I’ll argue that this assumption makes it difficult for Diaz-Leon to handle other kinds of cases and hence difficult to satisfy the usage data constraint compellingly.

3 Problems with Diaz-Leon’s Contextualism

In the previous section, I illustrated how Saul’s contextualism accommodates the usage data constraint. But I also illustrated how the view struggles to predict that, in Case 3, Ash is using ‘woman’ as a gender term. I then discussed how Diaz-Leon’s contextualism appears

¹³ See Schroeder (2007)

to improve Saul's, by accommodating the usage data constraint and predicting that Ash is using 'woman' as a gender term in Case 3.

Recall, however, that part of what motivates Saul to examine contextualism about 'woman' in the first place is its promise in accommodating cases in which speakers use 'woman' as a gender term *and* other cases in which speakers use it as a sex term. Diaz-Leon never circles back to discuss the latter. The main lesson of this section will be that Diaz-Leon's view has difficulty accounting for such cases. While Diaz-Leon's has the benefit of making the right prediction about Ash, it has the cost of failing to satisfy the usage data constraint. Or so I'll now argue.

Consider another example.

Case 4 Jo, who has XY chromosomes, penial genitalia, and the like, utters 'I'm a woman' while visiting a friend at a rape crisis center offering counseling to victims of rape, the overwhelming majority of whom at the center are female.

Mainstream usage of 'woman' suggests that speakers use it as a sex term in some scenarios. Case 4 is put forward as such a scenario – Jo is using 'woman' as a sex term. Suppose, however, that Jo is so fed up with being misgendered that they would become severely depressed and hence suicidal if 'woman' were not being used as a gender term.¹⁴ Plausibly, self-identification would then carry the most moral weight. After all, Jo's life in part depends on it. Because features that carry moral weight determine which standards are salient in context on Diaz-Leon's view, and because self-identification is a non-biological feature, Diaz-Leon's view predicts that a non-biological standard is salient in Case 4. Thus, Diaz-Leon's view predicts that Jo is using 'woman' as a gender term, not a sex term. That's a problem.

¹⁴ By using details involving a rape crisis center and depression-induced suicide in a philosophical example, I risk implicating that I do not take these issues seriously. Allow me to cancel that possible implicature: I take these morally significant issues deeply seriously.

This might strike readers as too quick, leaving them with the impression that *something* must be awry in Case 4. But Case 4 is no one-off counterexample. It is the result of a simple and repeatable recipe. Take any case in which ‘woman’ is plausibly being used as a sex term. Arrange the (consequentialist or non-consequentialist) features such that self-identification matters. Further arrange the features such that self-identifications matter more than any biological feature. The result will be that a non-biological standard is salient. Thus, speakers will use ‘woman’ as a gender term.¹⁵ Repeat these three steps.

Getting around this problem requires that self-identification carries enough moral weight to make a non-biological standard salient in cases like Case 3, but not so much moral weight as to make such a standard salient in cases like Case 4. That’s too fine a line to walk for Diaz-Leon’s view to make predictions that line up with the usage data. The point here is that, on Diaz-Leon’s view, ‘woman’ is used as a gender term in far more cases than is compatible with the usage data constraint. Diaz-Leon’s view sounds less like a vindication of the idea that ‘woman’ lives a double life than it does the idea that the term has, to abuse the metaphor, a hobby.¹⁶

In light of this line of criticism, it is natural to wonder whether Diaz-Leon’s view is a genuine form of contextualism rather than a form of *invariantism* in disguise, in which no part of the term shifts across contexts. Consider paradigmatically contextual terms like ‘I’ or ‘me’. We use these terms to talk about different speakers in different contexts. In fact, *every context with a different speaker* yields a different topic. Compare the number of contexts with different speakers to the number of contexts in which self-identification is not very morally significant. It’s no contest. But just as we would begin to doubt whether ‘I’ or ‘me’ are contextual if we only ever used them to talk about anyone besides *me*, the

¹⁵ It’s true that I am relying on the controversial idea (See Dowell 2016) that we can learn anything significant about the nature of *our words* by theorizing about the use of the graphemes that partially constitute them in various counterfactual scenarios. But this is a problem for everyone, since all the parties to this dispute employ the same sort of cases.

¹⁶ Over the course of developing my objection, I came to learn that Bettcher (2017) raises a similar worry with Diaz-Leon’s view. But Bettcher does so in the context of a survey article, and hence only has the space therein to gesture toward the problem. My presentation can be thought of as expanding on Bettcher’s compressed insight. Such convergence is encouraging.

author, we should doubt that ‘woman’ is contextual on Diaz-Leon’s view for similar reasons.

In reply, Diaz-Leon could insist that ‘woman’ is more analogous to terms like ‘I’ than it seems on a first pass. Sure, Diaz-Leon might say, because the world is such that self-identification matters morally, ‘woman’ is used as a gender term in a surprising number of cases. And sure, Diaz-Leon might also say, the world is such that ‘I’ is used in a huge variety of ways, e.g. as a *Jennifer-Saul-term*, *Esa-Diaz-Leon-term*, *Angela-Merkel-term*, *Serena-Williams-term* and so on. But, Diaz-Leon might then say, the source of this difference between ‘woman’ and ‘I’ is the nature of our *world*, not the nature of these *words*. In principle, the world could be such that we use ‘woman’ to talk about a huge variety of different standards in different contexts, perhaps equally as many standards as there are speakers occupying contexts. In other words, ‘woman’ and ‘I’ differ only superficially. Or so Diaz-Leon might say.

Grant the point that ‘woman’ and ‘I’ owe their differences to worldly features that can be ignored. Still, on Diaz-Leon’s view, there is a real difference between the nature of ‘woman’ and ‘I’, owing to the role that *normative* considerations play in uses of the former. Normative principles are widely held to be metaphysical necessary if true, e.g. necessarily there’s an obligation to treat persons with respect if there’s an obligation to treat persons with respect at all. If that’s right, then there will be a number of worlds that are such that e.g. misgendering amounts to a form of disrespect, and hence a number of worlds in which there’s an obligation not to misgender. In other words, there will be a number of worlds in which non-biological standards are salient, where ‘woman’ is used as a gender term. But then it seems like ‘woman’ is far less flexible than ‘I’. Thus, even granting the imagined response from Diaz-Leon, ‘woman’ seems different on their view from paradigmatically contextualist terms like ‘I’.

There is another response to these concerns with Diaz-Leon’s view. It’s no problem, one might claim, that Jo is using ‘woman’ as a gender term in Case 4. Yes, typical cases involving medical scenarios or rape crisis centers are ones in which biological features carry the most moral weight. But Case 4 isn’t typical – it’s a case where non-biological features

carry the most moral weight, in virtue of the fact that Jo would become strongly upset upon being misgendered. Accordingly, one might think, there's no problem with making an exception for Case 4.¹⁷

On one way of understanding this response, it's calling into question the data that constitutes the usage data constraint. Because any adequate characterization of 'woman' must accommodate the usage data involving the term, the usage data places constraints on an adequate characterization of 'woman'. The data at issue in this paper is that speakers sometimes use 'woman' as a sex term and sometimes use it as a gender term. Case 4 is put forward as an illustration of cases of the former type. The response above call this into question – it denies that Case 4 is one in which the speaker uses 'woman' as a sex term.

Keep in mind, however, that Case 4 is supposed to be illustrative. Even if it turns out to be a case in which the speaker is using 'woman' as a gender term, there will still be other, similar cases in which speakers are using 'woman' as a sex term where the non-biological features are such that Diaz-Leon's view predicts otherwise. Remember, Case 4 follows a simple and repeatable recipe. One could go through every such case, arguing that each of them are actually cases in which speakers are using 'woman' as a gender term. But doing so would seem to amount to a flat-out rejection of the usage data constraint, and the start of an entirely different debate.

In a last ditch effort, one might grant these worries but insist that they're not problems, because Diaz-Leo's view is put forward as a *revisionary* one. While it's true Diaz-Leon's contextualism makes predictions that are at odds with the usage data constraint, one might say, it's also true that speakers should put in some effort to use 'woman' in the way the view predicts we've been using it from here on.

Three comments: Firstly, Saul entertains the idea that we should consider contextualism as a revisionary proposal. (Saul 2012: 402). In response, Diaz-Leon is explicit about the nature of their project:

¹⁷ In correspondence, Diaz-Leon offers a reply in this spirit.

“In response, I want to argue that we can see this point [Saul’s claim that “disagreements over who counts as a woman are simply not to be settled by appeal to the facts of language”] also as a genuinely linguistic point, that is, as a *descriptive* claim about the (*actual*) meaning of woman, *not just as a revisionary proposal* about how the term woman should be used, given our moral and political aims...I want to argue that it is also possible to provide moral and political considerations that are relevant with regard to the *descriptive project* of finding out the meaning of woman.” (Diaz-Leon 2016: 248-249, emphasis mine)

Of course, however, while Diaz-Leon might not conceive of their project as revisionary, it nevertheless could be that it’s best conceived of as such. This brings us to the second response, which is more concessive. Yes, it very well could be that the prospects for contextualism are best when seen through the revisionary lens. This point is exactly right, and it’s one to which I’ll return toward the end of this paper. However, and this is the third response, it’s also worth noting that contextualism doesn’t, on the face of it, sit well with the revisionary goal of changing how ‘woman’ is used. For one of the central insights of contextualism is that speakers have less of a say, as it were, in how words are used, in virtue of the fact that *context* does a lot of the heavy lifting.

4 Toward an Alternative Account

Recall, I set out to determine whether contextualist views of ‘woman’ can satisfy the usage data communicative constraints. So far, however, I have said very little about the latter. This is in a way unsurprising. For one of the main selling points of contextualist views, generally speaking, is that they associate a *single* component with words across contexts, making the communicative constraint straightforwardly satisfiable. As such, the question taking center stage has not been whether contextualist views like Saul’s and Diaz-Leon’s can satisfy the communicative constraint. Rather, the question has been whether they can satisfy the usage data constraint. I’ve argued that Saul’s contextualism appears to do fairly well with the usage data constraint, but it has other problems accommodating cases like Case 3. Diaz-Leon’s contextualism accommodates Case 3, I’ve argued, but struggles to satisfy the usage data constraint.

In the space remaining, I briefly and tentatively explore the following idea: starting with views that have an easy time with the communicative constraint, like contextualism, might end up occluding a more promising route. Perhaps if we *start* with views that seem to have an easier time with the usage data constraint, we'll then end up with a view that satisfies it, the communicative constraint, and avoids troubles with accommodating Case 3.

From this perspective, the obvious alternative is the view that 'woman' is *ambiguous* – it has a sex-involving meaning *or* a gender-involving meaning, not a single meaning that determines different semantic values in different contexts. On such a view, speakers sometimes use 'woman' to communicate the sex-involving meaning but sometimes use it to communicate the gender-involving meaning. This view makes the usage data constraint a piece of cake.

But it also faces an awkward difficulty: the communicative constraint tends to motivate philosophers to *minimize* ambiguities whenever possible. Take a word like 'good' from ethics, for example. Rarely in everyday use do speakers have any trouble understanding what people mean with this term. This would be hard to square if 'good' were ambiguous, however. So, ethicists tend to conclude that 'good' and other ethically significant terms like it are not ambiguous.¹⁸

Recently, however, the idea that there are different *kinds* of ambiguities, some better than others with respect to preserving the appearance that speakers succeed often in communicating with one another, has been garnering attention in philosophy.¹⁹ A term that is ambiguous has multiple meanings. That much is familiar. Less familiarly, a term can be ambiguous in the sense that it is *homonymous*, having multiple *unrelated* meanings, or it can be ambiguous in the sense that it is *polysemous*, having multiple *related* meanings.

¹⁸ See Finlay (2014).

¹⁹ Polysemy has been receiving renewed attention in linguistics. See Vicente & Falkum (2017) for an overview.

Standard examples of terms that are homonymous include ‘bank’. This term has one meaning involving financial institutions and another involving rivers. Compare the homonymous ambiguity of ‘bank’ to the polysemous ambiguity of ‘cut’. The term ‘cut’ has a meaning involving incision, another involving turning sharply, passing objectionably in a que, insulting someone, and so on. The principle difference between ‘bank’ and ‘cut’ is that the meanings of the former appear completely unrelated to one another, while the meanings of the latter all seem to have something to do with the idea of division. Such relations appear to preserve our ability to communicate successfully – it’s much rarer to ask for clarification when an interlocutor uses a word like ‘cut’ than it is when they use a word like ‘bank’.

The point is that there need not be reasons to avoid an ambiguity approach to characterizing a term, and so there need not be reasons to avoid an ambiguity approach to ‘woman’.²⁰ That there might not be any communicative constraint-based reasons against embracing polysemous ambiguity for ‘woman’, of course, doesn’t imply that there are any reasons favoring one.

But it’s not hard to find such reasons. For example, there is little controversy surrounding the etymology of ‘woman’. Many would agree that large groups of speakers have only recently begun using it clearly as a gender term. This seems to imply that its recent meaning descends *historically* from an original one, which is interesting, because historical connectedness is one of the hallmarks of polysemy.²¹ Perhaps even more interestingly, a polysemous ambiguity account of ‘woman’ appears to have an important advantage over versions of contextualism like Diaz-Leon’s. Over the course of defending

²⁰ Leslie (2015) suggests a similar view explicitly, arguing in the context of the debate regarding generics, that it’s the best way to make sense of them.

²¹ See Viebahn & Vetter (2016)

contextualism, Diaz-Leon issues the following criticism of Bettcher (2013), who holds a related view²² to the one I am floating:

“...if we adopt a multiple-meaning account [or, relatedly, a polysemy account], we are accepting the existence of certain meanings that are exclusionary. That is, we will understand some speakers in some contexts as using *woman* in a way that is exclusionary...And this will have the consequence that whether (some) trans women count as women in some contexts will depend on which meanings are at issue...And this seems problematic.” (Diaz-Leon 2016: 254, original emphasis)

It’s true that there will be some meanings of ‘woman’ that are exclusionary on the view that the term is polysemous. However, and this is key, that’s exactly what should be expected on a realistic portrayal of the current socio-political state of the world. It’s hard to deny that mainstream usage of ‘woman’ often excludes trans individuals. That’s part of the reason there even exists trans-inclusive feminist activism. It’s precisely because so many speakers in the world use the excluding senses of ‘woman’ that trans-inclusive feminist activism has teeth. Trans-inclusive feminist activism would make little sense if everyone were already using ‘woman’ in a trans-inclusive way.

This way of understanding the significance of trans-inclusive feminist activism might seem to imply that philosophers should think about cases like Case 3 differently, where Ash uses ‘woman’ as a sex term in part because they are surrounded by traditionalists. Indeed, it does have this implication. Saul worries that in cases like Ash’s, the contextualist view they explore is morally problematic because it predicts that Ash is using ‘woman’ as a gender term. Partially as a result, Saul abandons the view. Diaz-Leon shares the worry, setting out to resolve it *semantically*, by developing a normative version of contextualism predicting that Ash is using ‘woman’ as a gender term. Saul and Diaz-Leon are right to

²² Bettcher (2013: 240-244) argues that ‘woman’ has both “dominant” and “resistant” meanings, which map, roughly, onto the polysemous view of ‘woman’ floated in this paper, where ‘woman’ has a meaning in mainstream society that is different from the meaning it has among trans and trans-inclusive feminist communities. This is why the views are similar. Though similar, I hesitate to characterize Bettcher’s view as polysemic not only because Bettcher doesn’t, but also because some of what Bettcher says suggests that speakers can “reject” as “false” attributions of terms with meanings for which they are against, implying that speakers have a degree of linguistic control that is in tension with my understanding of polysemic views. Comparing these views fully requires more space than is available in this paper.

worry that their views predict that Ash is using ‘woman’ as a sex term, but wrong to think that the remedy is to find a semantic view that doesn’t have this prediction. That’s a way of denying that Ash’s situation was really ever problematic, and hence a way of saying that there wasn’t to begin with any trans-inclusive feminist work to be done in such situations. But there was and is such work – it’s just more moral or political than semantic.

On a polysemic picture, one of the goals of trans-inclusive feminist activism is to agitate for broader use of the more inclusive senses of ‘woman’ available. To touch an earlier point, the resulting picture is a revisionary or engineering one in a particular sense. The goal isn’t to conceptually engineer a new concept, in the sense of introducing one. That’s already happened, as use of it among trans and trans-inclusive feminist communities indicate. Instead, at least from the perspective of trans-inclusive feminist activism, the goal is to make the inclusive use of ‘woman’ more mainstream.²³

It might be said in response that the polysemy view doesn’t have the advantage of making better sense of trans-inclusive feminist activism than Diaz-Leon’s contextualism. Suppose Diaz-Leon were to supplement their brand of contextualism with resources broadly available to semantic externalists, on which words, roughly, refer to what they do in virtue of being causally connected via chains of use to baptismal events introducing them or in virtue of being connected deferentially to use among communities of experts. Details aside, the upshot could be that mainstream use of ‘woman’ is maximally inclusive but speakers aren’t aware of it.²⁴ If so, then it could be that one of the organizing goals of trans-inclusive feminist activism is to make everyone aware that they’re using the word ‘woman’ inclusively. Perhaps that’s the best way to understand the significance of trans-inclusive feminist activism, instead.

This response raises a number of interesting issues that there isn’t space to explore fully. Nevertheless, there are at least two reasons – one specific to the debate and one more general – to think it doesn’t succeed in showing that Diaz-Leon’s view of ‘woman’ makes

²³ Dembroff (forthcoming) arrives at conclusions in the same spirit from a different starting point.

²⁴ See Kripke (1972) and Putnam (1975) for canonical treatments. See Mallon (2017) for a recent discussion of externalist views of social terms, errors in use, and other related issues.

better sense of trans-inclusive feminist activism surrounding it than a polysemous view. Suppose speakers are using ‘woman’ in a maximally inclusive way on a version of Diaz-Leon’s view that is supplemented with externalist resources. If so, and this is the first specific reason, then it’s hard to see why Saul and Diaz-Leon would ever have been worried about Ash’s use of ‘woman’ in cases like Case 3. But even if they’re mistaken to worry, as I’ve suggested above, it’s not hard to see why they’re worried.

Moreover, and this is the second more general reason, the English language is *shot through* with words that are polysemic.²⁵ Though the issue is seriously underexplored, the fact that the English language, and most languages, for that matter, contain so many polysemic words seems to be in serious tension with strong forms of semantic externalism. Combining strong forms of semantic externalism with the fact that the English language contains wildly many polysemic words seems to suggest an implausible degree to which English speakers are ignorant or in error.

5 Conclusions

One of my main aims in this paper has been to show that Diaz-Leon’s version of contextualism predicts that we use ‘woman’ *much* more often as a gender term than as a sex term, in a way that is starkly at odds with the usage data constraint. Along the way, I also put forward the claim that we should be open to ambiguity-based alternatives to contextualism, like a polysemy account of ‘woman’. Admittedly, it’s easy to walk away from these claims with the impression that the central upshot of this paper is that the whole family of contextualist views about ‘woman’ is hopeless.

But that’s not quite right. As it happens, the jury is still very much out on how to understand the precise differences between contextualism and polysemy. Indeed, some philosophers, like Recanati (2017), argue that there isn’t much daylight between polysemy

²⁵ As Fogal (2016) points out, “...polysemy is utterly pervasive in natural language...and it affects both content and function words. It also tends to be both systematic and productive, with similar patterns of polysemy applying to similar words across many languages...the default hypothesis for pretty much any ordinary noun, adjective, or verb should be that it is polysemous, and thoroughly so.”

and contextualism. It could be, then, that in defending a polysemic view of ‘woman’, I’ve defended a kind of contextualism. Moreover, I argued that attempts like Diaz-Leon’s to make facts about the contextualist semantics of ‘woman’ line up with facts about how we ought to treat each other morally, *by building morality into semantics*, come up short. But not all contextualist views build morality into semantics. The contextualist view that Saul floats doesn’t. So, I can also be read as defending *non-normative* versions of contextualism like Saul’s; at least, I can be read as offering such a defense, so long as what’s really going on in cases like Ash’s in Case 3 is appreciated – their context is such that ‘woman’ is being used as a sex term now, but trans-inclusive feminist activists are working to create a society in which that’s not the case later.

Whether a non-normative contextualism or polysemy ambiguity theory for ‘woman’ is best, it’s worth emphasizing that such views are promising in part because they don’t incorporate moral conditions into semantics. Instead of tasking semanticists with sorting out how it could be that the semantic facts line up with the moral facts, we should ask moral philosophers, including, especially, trans-inclusive feminist moral philosophers, what we ought to do in light of these facts failing to line up.

References

- Barnes, E. (2019). “Gender and Gender Terms”. *Nous*. <https://doi.org/10.1111/nous.12279>
- Bettcher, T. 2013. “Trans Women and the Meaning of “woman.”” In *The Philosophy of Sex*, ed. Nicholas Power, Raja Halwani, and Alan Soble. Lanham, Md.: Rowman and Littlefield.
- Bettcher, T. 2017. “Trans Feminism: Recent Philosophical Developments”. *Philosophy Compass* 12 (11):e12438.
- Dembroff, R. & Wodak, D. 2018. “He/She/They/Ze”. *Ergo: An Open Access Journal of Philosophy*.

- Dembroff, R. (forthcoming). “Real Talk on the Metaphysics of Gender”. *Philosophical Topics*.
- Diaz-Leon, E. 2016. “Woman as a Politically Significant Term: A Solution to the Puzzle”. *Hypatia* 31 (2):245-258.
- Dowell, J. 2013. “Flexible Contextualism about Deontic Modals: A Puzzle about Information-sensitivity”. *Inquiry : An Interdisciplinary Journal of Philosophy* 56 (2-3):149-178.
- Dowell, J. 2016. “The Metaethical Insignificance of Moral Twin Earth”. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics*. Oxford University Press.
- Finlay, S. 2014. *Confusion of Tongues: A Theory of Normativity*. Oxford: Oxford University Press.
- Fogal, D. 2016. “Reasons, Reason, and Context”. In Errol Lord & Barry Maguire (eds.), *Weighing Reasons*. Oxford University Press.
- Haslanger, S. 2000. *Gender and Race: (What) are they? (What) do we want them to be?* *Nous* 34 (1): 31–55.
- Jenkins, K. 2016. “Amelioration and Inclusion: Gender Identity and the Concept of Woman”. 126 (2): 394–421.
- Kripke, S. (1972/1980). *Naming and Necessity*. Cambridge, Massachusetts, Harvard University Press.
- Leslie, S. 2015. “Hillary Clinton is the Only Man in the Obama Administration”: Dual Character Concepts, Generics, and Gender. *Analytic Philosophy* 56 (2):111-141.
- Mallon, R. (2017). “Social Construction and Achieving Reference”. *Noûs* 51 (1):113-131.
- Mikkola, M. 2017. “Feminist Perspectives on Sex and Gender”. In *The Stanford Encyclopedia of Philosophy*, ed. Edward Zalta. October 2017 edition. <http://plato.stanford.edu/>
- Putnam, H. (1975). “The Meaning of “Meaning.”” *Mind, Language and Reality: Philosophical Papers*. New York, Cambridge University Press. 2: 215–271.

- Recanati, F. 2017. "Contextualism and Polysemy". *Dialectica* 71 (3):379-397.
- Saul, J. 2012. "Politically Significant Terms and Philosophy of Language: Methodological Issues. In *Out from the Shadows: Analytical Feminist Contributions to Traditional Philosophy*, ed. Sharon Crasnow and Anita Superson. Oxford: Oxford University Press.
- Schroeder, M. 2007. *Slaves of the Passions*. Oxford: Oxford University Press.
- Vicente, A. & Falkum, I.L. (2017) "Polysemy". In *Oxford Research Encyclopedia of Linguistics*. Ed. Mark Aronoff. New York: Oxford University Press.
- Stephanie, K. 2016. "Misgendering and its Moral Contestability". *Hypatia: A Journal of Feminist Philosophy* 31 (3):512-519.
- Viebahn, E. & Vetter, B. 2016. "How Many Meanings for 'may'? The Case for Modal Polysemy". *Philosophers' Imprint* 16.