

Not strictly a woman – QUD-based 4-valent reasoning discharges lexical meaning*

Emil Eva Rosina

Ruhr University Bochum, Bochum, Germany
emil.rosina@ruhr-uni-bochum.de

Abstract. I offer a framework that captures both context-dependency and vagueness of predicate meanings – illustrated by the politically relevant case of *woman* – as an interaction of lexical meaning and Question under Discussion (‘QUD’). I extend [8]’s non-maximality approach to superficially polysemous predicates like *woman* and show that this is conceptually coherent and insightful for a linguistic analysis of political debates about gender invitation policies: While there are (i) clear, semantically true, and (ii) strictly false cases of *x is a woman*, there are also (iii) merely pragmatically acceptable cases (‘like a woman wrt. the QUD’) as well as (iv) truly vague ones. I argue that this four-way division is the least complex model that captures current gender discourses in a harm-reducing, trans-inclusive way. This offers a new perspective on the semantics-pragmatics interface of predicate meanings.

Keywords: Gender predicates · Question under Discussion · Polysemes.

1 Introduction

1.1 Starting Point

This paper introduces a new framework for trans-inclusionary usage of gender predicates like *woman* as well as a new take on polysemy. I focus on the following case where no obvious choice of lexical meaning¹ of the term *woman* is acceptable in the light of the boldfaced ethical intuitions in (1-b)–(1-f) and the goal to not misgender people:²

* This paper is based on a joint talk with Franci Mangraviti at the Internal Workshop of the Vienna Forum for Analytic Philosophy on the 20th of February 2023. I owe both motivation and background for my linguistic work in this paper to political and philosophical discussions with Franci Mangraviti – thank you!! I also thank Jordan Chark, Peter Sutton, Nina Haslinger, Kristina Liefke, the audience of the Internal Workshop of the Vienna Forum for Analytic Philosophy as well as three anonymous ESSLLI reviewers for their helpful comments. The research for this paper is supported by the German Research Foundation, DFG, as part of Kristina Liefke and Markus Werning’s project in the research unit FOR 2812: Constructing Scenarios of the Past (grant nr. 397530566). Beyond its main scope, I also want this paper to be understood as a commitment to integrate my political and academic believes.

¹ I am not making any assumptions about the nature of this choice/“choice”.

² I take misgendering to be equivalent with not respecting someone’s gender self-ID.

- (1) a. SCENARIO GLASS CEILING: An agency provides funding for a ‘Fight the glass ceiling in academia’ event with the goal of providing tools to counteract gender-related imposter syndrome, gender CV bias, etc. According to the invitation, the event is ‘for women only’.
- b. Trans woman Tammy wants to join. (x =Tammy) **invite!**
- c. Non-binary person Noa (mostly read female, suffering from imposter syndrome) wants to join. (x =Noa) **invite!**
- d. Trans man Mario (transitioned as a teenager, always read male) wants to join. (x =Mario) **invite?**
- e. Assigned male at birth Quinn (recently started questioning their gender, always read male and masculine presenting) wants to join. (x =Quinn) **invite?**
- f. Cis man Chris who will never question his gender and is perceived as masculine and determinate wants to join. (x =Chris) **don’t invite**
- g. $\llbracket x \text{ is a woman} \rrbracket = 1 \text{ iff } \dots ?!$

I will discuss several options for choosing the lexical meaning of *woman* in relation to this case and conclude that none of them does the job. Since my illustration will keep the scenario constant, the solution cannot be ordinary polysemy resolution in dependence of the context either. I argue that independently-motivated linguistic tools can be used to model good practice in this case³ and propose that a certain interaction between semantic truth-conditions, pragmatic acceptability-conditions, and vagueness captures the ethical intuitions above.

1.2 Preview and structure

Section 2 investigates trans woman Tammy’s (1-b) and non-binary Noa’s (1-c) position in the GLASS CEILING scenario. Subsection 2.1 discusses the political problem, subsection 2.2 provides the relevant linguistic tools, subsection 2.3 brings the two together and puts forth a first proposal for the meaning of *woman* that captures the intuitive ethical judgement regarding the cases of Tammy and Noa. At the end of Section 2, the picture developed will include three (out of a total presented four) truth-/acceptability-values: semantic truth, pragmatic acceptability, and falsity. Tammy (1-b) is invited qua the semantic truth of *x is a woman* where *woman* = ‘*x* self-identifies as a woman’. Noa (1-c) is invited on the basis of pragmatic acceptability, i.e. being *like* a woman in the respects relevant for the GLASS CEILING scenario which is modelled via answer-equivalence relative to the Question under Discussion (‘QUD’).

³ I do not claim that my solution is how most people actually use the term *woman* in current English. If, however, our political goal was impossible to obtain on linguistic grounds, it would be subject to a ‘*ought* implies *can*’ objection. I propose a use of *woman* that may currently be used by some speakers but might in principle be acquired by everyone, because QUD and vagueness are elements of natural languages anyway. At the same time, I point to a set of possible tools for modelling (what is wrong with) non-optimal actual uses of gender terminology, see section 4.1.

Section 3 focuses on the case of trans man Mario and questioning Quinn. Again, the first subsection, 3.1, makes the political case, the second one, 3.2, provides the linguistic background, and the third subsection, 3.3, brings them together for a solution. Mario (1-d) and Quinn (1-e) are modelled as truly vague (or ‘tolerantly acceptable’ cases), such that their in- or exclusion depends on further, possibly unresolvable, factors. At the same time, the concept of tolerant acceptability serves to stop any Sorites paradox (introduced in subsections 3.1, 3.2) that might seem to lead to inviting Chris as well. At the end of Section 3, all four truth/acceptability values have been introduced and the basic idea is complete. A table sums up how I model all of the intuitions in (1).

Finally, Section 4 addresses three loose ends: Subsection 4.1 shows that my account can also be used to model bad practice. Subsection 4.2 contains some pointers to a fruitful probabilistic extension. In subsection 4.3, I conclude with possible implications of my account at the semantics-pragmatics interface.

1.3 Methodological remark

This paper is methodologically non-standard both from a gender-philosophical as well as from a linguistic point of view. I am applying a framework that has been developed for non-maximal determiner-phrases (like *the windows*, as introduced in Section 2.2) to a what seems to be a polysemy problem. Existing accounts to polysemy didn’t seem to be able to model how I *would like to* use gender terms like *woman*. If I hadn’t considered that a problem, I wouldn’t have had any motivation to point to the parallel, because existing accounts of polysemy are able to deal with the politically less loaded cases. That said, the parallels between context dependency in gender philosophy and as the technical term used in linguistics should not be ignored, I argue, if intuitions map as neatly as in this case – and my claims do not have to be restricted to the politically relevant cases, see Subsection 4.3. So while a politically normative motivation was the starting point of this project, it has turned into one of several models for describing the use of seemingly polysemous predicates in general. The normative power carries on to choosing this model over others that don’t make sense of the politically relevant cases. As for the original political motivation, I am confident that the account presented here provides the tools for better language policies, but I do not illustrate how to put it into practice (e.g. how to actually phrase an invitation or how to decide in the vague cases in practice). I leave this as well as the relation to philosophical gender models to future gender-philosophical work.

2 *Like* a woman – Truth and QUD-based acceptability

2.1 Political Background: Gender Invitation Policies

Events that used to invite only ‘women’ (e.g. to counteract male supremacy) pose an ethical problem that relates to word-usage: Descriptively, people have acquired different concepts of gender that influence their choice of words. Most of

these misgender some group of people, e.g. whether one uses $woman_2$ =‘assigned female at birth’, $woman_3$ =‘low on the patriarchal hierarchy’, $woman_4$ =‘has a mainly estrogen based body’, $woman_5$ =‘has an F in their passport’ or any combination of these, one will always misgender someone, except for $woman_1$ =‘self-identifies as a woman’.⁴ Since, for many people, their gender self-identification does not change in dependence of the context,⁵ and given that we want to reduce misgendering harm, it seems that $woman_1$ should be(come) the one and only lexical meaning of *woman* without allowing for any contextual flexibility:

(2) $\llbracket x \text{ is a woman} \rrbracket = 1$ iff x self-identifies as a woman

However, using $woman_1$ in contexts like the one described by (1-a), GLASS CEILING, (that target gender-related disadvantages) will exclude people who are subject to the target discrimination, the clearest case being ‘visibly’ trans masculine people like Noa in (1-c). At least, one might say, the organizers are not so trans-exclusive as to apply some other notion of *woman* (like $woman_2$ =‘assigned female at birth’), and so Tammy (1-b) is able to join the event (independent of her clothing, body features etc.). At the same time, something seems wrong, if Noa (1-c) is excluded even though they are subject to the kind of discrimination which the event targets. Noa’s position is doomed: Trying to include them while using *woman* on the invite does misgendering harm (because they would be invited as a woman without identifying as one), excluding them harms them via withholding tools of empowerment. Note that the problem is not just that there is no consistent way *across* contexts of using *woman* without doing harm – a problem which could be solved by many versions of contextualism. Rather, there is not even a good choice of words *within* one scenario like GLASS CEILING, which is a strong objection to standard polysemy accounts: Using $woman_3$ =‘low on the patriarchal hierarchy’ in this context invites both Tammy and Noa in.⁶ However, in Noa’s case, this happens on the basis of calling them a woman. Meanwhile, many other options $woman_n$ draw on irrelevant factors and just render more combinations of exclusion and misgendering harm.

A salient answer to this problem is to never use the term *woman* on invitations of this kind (but something like *people who are low in the patriarchal hierarchy*). This is however open to attack with arguments from practicability and classism: If we are not understood, stereotypically in rural working-class contexts, and no-one joins our event, there is no use in a theoretically good choice of words.

⁴ I am treating self-identification as a black-box. Many people do report causal relations between self-ID and some of the other aspects (like body schema, clothing, a.o.) often associated with gender. However, neither of them is a reliable predictor of self-ID. I am not making any claims about the metaphysics of gender here.

⁵ I am aware that gender-fluidity contrasted with retrospective self-ID complicates things. Nonetheless, I have to exclude the time component for matters of space.

⁶ Yes, trans women are low on the patriarchal hierarchy. See also sections 4.1 and 4.2 for the importance of this implication.

If we want to stick to using *woman* in some contexts, then, modelling the normatively best option – inviting everyone who identifies as a woman, plus everyone who is subject to the disadvantages that are to be counteracted for different reasons (for philosophical motivation see e.g. [11]) – linguistically involves some sort of pragmatic slack or mistake: We are looking for a way of including Noa without calling them a woman, while at the same time using the term *woman* on the invite. I will show that the paradoxical impression of this endeavour is merely a result of focusing on strict semantic truth and falsity, while the relevant cases involve a specific kind of context-dependency (this Section 2) and vagueness (Section 3) and require a minimum of four truth/acceptability values.

2.2 Linguistic Background: Question under Discussion

I point to a parallel between a problem that has been discussed with respect to non-maximal meanings (as in (3-b) below) of sentences involving definite determiners (*the*) and cases like GLASS CEILING. While definite DPs like *the windows* in (3-b) give rise to the same (semantic) truth-conditions as universal quantifier phrases (*all windows*) in sentences like (3-b), these are often accepted in non-maximal scenarios like THUNDERSTORM.

- (3) a. SCENARIO THUNDERSTORM: Ann and Bea left their flat. A thunderstorm is coming. 4 of the 10 windows are open. Ann: *Oh no!...*
 b. *The windows are open!* **QUD:** Will our flat be flooded?

The above has been explained via what I will be calling (pragmatic) acceptability conditions that can be much weaker than the (semantic) truth-conditions (‘literal meaning’) of a sentence, depending on features of the context. E.g. [12], [13], building on [15], resort to the concept of Question under Discussion (‘QUD’) – roughly: the motivation for the current move in a conversation; the question whose answer requires speakers to know the status of the windows:⁷ The salient QUD in THUNDERSTORM is something like ‘Will our flat be flooded?’ The answer is plausibly the same (‘yes’) regardless of whether 10/10 or 4/10 windows are open, i.e. these two states of affairs are in the same partition with respect to this QUD, 4/10 windows is *like* 10/10 windows for the flooding QUD. (3-b) is merely pragmatically acceptable in all cases where it is not semantically true, but this is irrelevant for the purpose of answering the QUD. Definite determiners license such non-maximal construals while universal quantifiers do not. That context-dependency in the form of QUD-dependency is a plausible root of the acceptability of (3-b) in THUNDERSTORM becomes clear in contrast with (4):

- (4) a. SCENARIO PAINT: Ann and Bea left their newly painted flat. The paint requires maximal ventilation. 4 of 10 windows are open. Bea: *Did you leave everything like the instructions say?* Ann: *Yes! ...*
 b. *??The windows are open!* **QUD:** Will the paint dry?

⁷ See e.g. [5] for a more complete picture.

Ann’s answer seems at least dubious in this case, since the new scenario gives rise to the QUD ‘Will the paint dry?’, for which the difference between 10/10 and 4/10 open windows does matter, i.e. these cases are *not* in the same partition with respect to this QUD.

2.3 Application: QUD-dependence in the GLASS CEILING case

Extending QUD-based acceptability conditions to cases with gender predicates solves language policy problems. I suggest that the more general equivalent of a non-maximal scenario is one that doesn’t make a sentence semantically true, but pragmatically acceptable because of QUD-answer-equivalence to the semantical truth evoking case. Let’s reconsider GLASS CEILING in this light:

- (5)
- a. GLASS CEILING, **QUD**: Is x low on the patriarchal hierarchy?
 - b. Trans woman Tammy wants to join. (x =Tammy)
 - c. Non-binary person Noa (mostly read female, suffering from imposter syndrome) wants to join. (x =Noa)
 - d. *x is a woman.* **true for Tammy (5-b), acceptable for Noa (5-c)**

In order to include Tammy (regardless of her passing, gender performance and body features), we want to use $woman_1$ =‘self-ID woman’ as the literal meaning, so (5-d) is semantically true in subscenario (5-b). Adding acceptability via QUD-sensitivity allows Noa in as well without (clearly) misgendering them and makes (5-d) acceptable in subscenario (5-c). Noa is *like* a woman with respect to this specific QUD, but unlike one in many other respects. Importantly, it is not even pragmatically acceptable to say *Noa is a woman* just because we are in the broader context of the event, because the QUD is more fine-grained. Being low on the patriarchal hierarchy is just an appropriate QUD when we decide who to invite, talking about a group. Any out-of-the-blue utterance in no matter which environment which addresses Noa as an individual is arguably ‘about’ their gender (i.e. self-ID) in a direct sense.⁸

The central move here, in relation to the dilemma sketched in subsection 2.1, is that only one of the aspects of the seemingly polysemous lexical entry of *woman* actually constitutes the lexical entry, namely $woman_1$ =‘self-identifies as a woman’. The other relevant factor is untied from $woman_3$ =‘low on the

⁸ I thank an anonymous reviewer for pushing me to comment on this point with regard to Quinn and Mario (subsection 3.3). I take this issue to be even more pressing in the strictly pragmatically acceptable case of Noa. Being a trans-masculine non-binary person myself, I would not feel misgendered by the sort of pragmatic acceptability involved here, if I were Noa. On the contrary, I embrace the idea of being in the same relation to a woman as are the 4 open windows to the 10 open windows: They are just really not the same, but the circumstances cluster them together. That said, I cannot speak for everyone belonging to this group. But remember also that we are dealing with an ethical dilemma that might not be resolvable without doing any harm at all. Much practical philosophical work has to be done in weighting misgendering, exclusion and unaccessibility/classism harm in cases like GLASS CEILING.

patriarchal hierarchy’ and serves as a QUD now that is only indirectly related to gender (in that every or practically – see subsection 4.2 – every $woman_1$ is low on the patriarchal hierarchy in the actual world).⁹

3 Vaguely like a woman – Tolerant acceptability

3.1 Political Background: Slippery Slope-ism

Organizers of events like GLASS CEILING who try to avoid misgendering and exclusion of people who are discriminated against from valuable resources and experiences are subject to slippery slope arguments that are used to justify trans-exclusive practices. Usually, their proponents voice a variation of the following worry: ‘If we invite trans women regardless of how they dress and how their body looks, and/or if we allow assigned-male-at-birth people in who have just started questioning their gender a bit, we might as well let cis men in as well, because we cannot tell the difference and/or because the boundaries are fuzzy.’

Since I already established $woman_1$ =‘self-ID woman’ as the semantic truth criterion, this argument does not go through for Tammy even if she doesn’t take hormones and is butch presenting.¹⁰ There are, however, cases where even trans-inclusive intuitions about whether to invite someone to the GLASS CEILING event are unclear, and we don’t want that to lead to inviting everyone in:

- (6)
- a. SCENARIO GLASS CEILING
 - b. Trans man Mario (transitioned as a teenager, always read male) wants to join. (x =Mario) **invite?**
 - c. Assigned male at birth Quinn (recently started questioning their gender, always read male and masculine presenting) wants to join. (x =Quinn) **invite?**
 - d. Cis man Chris who will never question his gender and is perceived as masculine and determinate wants to join. (x =Chris) **don’t invite**

Depending on further factors (like Mario’s possible imposter syndrome resulting from his childhood socialization, his level of toxic masculinity in academia, the ‘strength’ and nature of Quinn’s questioning, ...) that are hard or impossible to determine, the QUD ‘Is x subject to gender based discrimination (to a significant degree)?’ or ‘Is x low on the patriarchal hierarchy?’ might be answered positively or negatively for Mario and Quinn, and it seems impossible to pinpoint to the exact degree of change in their psychology or treatment by their environment that would determine the answer. If we act permissively in these unclear cases, so the worry, there is no clear ‘degree of cis-masculinity’ that could serve as a flipping point, and so there is no clear reason for not letting Chris join as well.

⁹ This can be linked to the metaphysical relation between gender self-ID and patriarchal oppression: Are women oppressed by virtue of being women? see e.g. [7],[6]

¹⁰ Instead, it comes down to not believing people what they say is their gender.

3.2 Linguistic Background: Vagueness

Things get more complicated than the three-fold distinction ‘semantic truth – mere pragmatic acceptability – falsity’ once we look at cases like (7), a variation of scenario PAINT with 70/100 open windows:

- (7) a. SCENARIO PAINT 2: Ann and Bea left their newly painted luxury estate. The paint requires maximal ventilation in order to dry. 70 of the 100 windows of their mansion are open. B: *Did you leave everything like the instructions say?* A: *Yes!...*
 b. *?The windows are open!* **QUD:** Will the paint dry?

While semantic truth requires all windows to be open and no open windows makes (7-b) clearly false, 99/100 windows constitutes an acceptable case here relative to the QUD ‘Will the paint dry?’. However, intuitions on (7-b) in SCENARIO PAINT 2 with 70/100 open windows might vary and lead to unclear judgements. This is strongly reminiscent of standard cases of vagueness like with ‘x is tall’ where (i) there are borderline cases and (ii) that give rise to a version of the Sorites paradox (see e.g. [9]). Sorites reasoning can also be applied to opening one window after the other: A particular speaker judges (7-b) false in SCENARIO PAINT 2, and they agree that one of a total 100 windows never changes ventilation enough to make a difference. Then, it seems, they are committed to rejecting a version of SCENARIO PAINT 2 with 71 open windows, and consequently one with 72 open windows, and so on, until they have to disagree that 100 open windows are enough for the paint to dry, which is highly counter-intuitive.

Treating vague cases like (7-b) in SCENARIO PAINT 2 as tolerantly acceptable (the forth ‘truth/acceptability value’ I am going to use)¹¹ stops the modus ponens that enables the Sorites paradox: One of 100 windows does not make a difference between strict acceptability and falsity, but it might make a difference between strict acceptability and tolerant acceptability, or between tolerant acceptability and falsity. Following [8], building on [1], tolerant acceptability arises when the combination of QUD and broader context involves vagueness, i.e. when the explicit context justifies several different interpretations of what ‘counts as’ enough open windows / dry enough paint, and consequently of which cases are in the same partition as the strictly true/acceptable cases. What I have been calling ‘(pragmatic) acceptability’ in Section 2 must now be more precisely termed ‘strict (pragmatic) acceptability’ in contrast to ‘tolerant (pragmatic) acceptability’, i.e. ‘real’ vagueness (that is not resolved by the QUD).

The tolerant cases are the ones that ‘are like’ the strictly true/acceptable ones with respect to some partition of the QUD, but not with respect to another partition. Applied to (7-b) in SCENARIO PAINT 2, this means that there is at least one reasonable interpretation under which 70 of 100 windows are enough for the paint to dry sufficiently for the purpose at hand (and so 70/100 is pragmatically acceptable because of QUD-answer equivalence to 100/100 windows). But there is also at least one reasonable interpretation under which the paint

¹¹ I am not making any ontological or logical claims, this is just intuitive terminology.

will not dry sufficiently with 70/100 windows open (and so this case partitions with the clearly false cases, like 0/100 or 2/100 windows open, and is not even pragmatically acceptable). The way [8] views it, this is the additional complication of ‘real’ vagueness as opposed to the kind of context dependency that leads to clear pragmatic (un-)acceptability once the QUD is fixed.

3.3 Application: Tolerant acceptability in the GLASS CEILING case

[8]’s framework leaves room for unresolved vagueness (tolerant acceptability) AND clear-cut cases of context-dependency (strict acceptability) while being immune to slippery slope arguments resembling the Sorites paradox: there are also strictly true as well as clearly false cases. This makes it a perfect shield against the trans-exclusive objection that acknowledging borderline cases and fuzzy boundaries / treating these cases permissively will also allow cis men in:

- (8) a. GLASS CEILING, **QUD**: Is x low on the patriarchal hierarchy?
 b. x is a woman. **tolerantly acceptable for Mario (6-b) and Quinn (6-c), false for Chris (6-d)**

It is adequate to not resolve the vagueness in these cases (although it may be worth debating whether it can be reduced to epistemic shortcomings and under-specification) – and in practice to let Mario and Quinn decide for themselves in most cases. The resistance to give a clear positive or negative judgement arises exactly in those cases where the underlying issue is too vague itself / gives rise to multiple interpretations in the given context to even judge pragmatic acceptability. This is reminiscent of the 70/100 windows case: For both Mario and Quinn, there is at least one reasonable interpretation under which they are sufficiently low on the hierarchy of patriarchy for the purpose at hand for x is a woman to count as pragmatically acceptable because of QUD-answer equivalence to women.¹² At the same time, there is also at least one reasonable interpretation under which Mario and Quinn pattern with cis men with respect to the QUD.¹³

¹² In Mario’s case one might argue the placement in the hierarchy of patriarchy may be carved into his thinking and feeling long after all visible signs of assigned at birth sex have vanished. In Quinn’s case, the mere questioning of their gender might already put them in a vulnerable position in a dichotomizing society to a degree that affects their social work dynamics and privileges substantially.

¹³ Mario’s self-socialization as a man could’ve lead him to behaving in toxically masculine ways at work, plus everyone could be treating him exactly like a cis man because no-one knows of his being trans. Quinn’s questioning might not have reached any emotional or even personal level but remained a theoretical ‘gender is a construct’.

It’s open to Mario and Quinn to contribute to resolving the vagueness further, if possible, by specifying for themselves e.g. what it means for Mario to be a man, and what the nature of Quinn’s question is for them. That would equal the following extra information in SCENARIO PAINT 2: ‘This flyer says that in case you have 100 windows, at least 90 of them have to be open for the paint to dry sufficiently to not leave stains on clothes.’ (But there might also be aspects of this vagueness that are not even *in principle* resolvable.)

Mario and Quinn’s status does not imply inviting clear cis men like Chris (6-d) to our event.¹⁴

Table 1. Truth/acceptability conditions of *x is a woman* in scen. GLASS CEILING

	iff...	Tammy	Noa	Mario	Quinn	Chris
true	<i>x</i> self-identifies as a woman	✓	×	×	?	×
strictly accept.	<i>x</i> is in the same partition wrt. the QUD as <i>y</i> , and <i>y</i> self-identifies as a woman ¹⁵	✓	✓	?	?	×
tolerantly accept.	<i>x</i> is in the women-partition wrt. some interpretation of the QUD but not wrt. some other	×	×	✓	✓	×
false	neither of the above	×	×	×	×	✓
invited?		✓	✓	?	?	×

4 Open issues

4.1 Modelling ethically bad use of gender terms

My four-way QUD-based approach can also model bad practice: A group of ‘trans exclusionary radical feminists’ (‘TERFs’) argues that granting trans people certain rights threatens the progress that has been made for cis women. Many ‘TERFs’ don’t consider trans women ‘real women’, e.g. using exclusively *woman*₂=‘assigned female at birth’ – not “only” misgendering people, but also contributing to a basis for trans hate with sometimes deadly consequences. ‘TERF’ events are typically restricted to all and only assigned-female-at-birth people who are the exact extension of *women* as used by ‘TERFs’. Alternatively, ‘TERFs’ may ‘accept’ the existence of trans women terminologically, but reject that they are subject to the same or similar oppression as cis women.

An application of the four-valued framework is that it captures many ‘TERF’s’ word use as missing context-dependency by conflating truth and acceptability: For many ‘TERFs’, being a woman is an in-or-out-criterion across contexts, not only in the semantic sense, but also with respect to acceptability. This misses the fact that different events that were formerly dedicated to ‘women’ in binary opposition to ‘men’ serve different purposes, so different QUDs are adequate, and the pragmatic extension of *woman* must therefore vary.¹⁶

¹⁴ ...although Chris is an artificially clear case and the exclusion of any individual is arguably often not the best solution to structural problems, which is especially relevant in similar cases with the goal of providing maximal safety at a party.

¹⁵ This is equivalent to being in the same partition as *all* women, as long as the implication from ‘woman’ to ‘low in the patriarchal hierarchy’ holds, see section 4.2.

¹⁶ This is by far not the only fallacy in ‘TERF’ reasoning, of course: Their choice of lexical meaning does misgendering harm. ‘TERF’s’ believe that trans women are

4.2 A probabilistic extension?

If ‘TERFs’ who believe that many trans women are not low on the hierarchy of patriarchy were right, QUDs about patriarchy/discrimination would not fit the literal meaning $x \text{ is a woman}_1 = \text{‘self-ID woman’}$, because then, being a woman would not imply being affected by the reason for the event.¹⁷ However, the way I put this so far, it may seem like a single woman₁ who is not *de facto* discriminated against at all at work could destroy the whole line of argumentation. While I am not aware of such a case, this would be a sore point, if a certain idealization (and, consequently, no strict implication in practice) wasn’t also present in the standard windows cases: Not in every possible world do 10/10 open windows during a thunderstorm have the consequence of a flooded flat (the wind might be going in the other direction). Playing around with the degree to which this implication has to hold is a promising project for future work. E.g. strict truth could be required to be a reliable predictor for the answer to the QUD for the QUD-literal combination to be admissible. I.e. how likely it is that if you are a woman₁ you are subject to gender discrimination might be relevant.¹⁸

4.3 Consequences for the pragmatics-semantics interface?

I have made a case about a noun predicate of the type $\langle e, t \rangle$ that is superficially (i.e. pre-theoretically) polysemous and where the division into four different truth/acceptability values is particularly necessary. Once we take the costs of extending the domain of the four-valent approach from non-maximality to polysemy in principle, however, there is no reason why politically loaded predicates should be special in this respect: Whether someone counts as a grown-up with respect to who gets to ride the roller-coaster seems to give rise to pragmatic acceptability (not a grown up, but as tall as one), while in the context of who gets to sit at the grown-up table, the concrete flipping point is unclear, so we get tolerant cases.¹⁹ Nor does the syntactic category intuitively play a role, when I apply the same reasoning to the adjective *female*. Maybe this phenomenon is restricted to the predicate type $\langle e, t \rangle$, then, or maybe it is not restricted at all:

not suffering from patriarchy (enough / as much as cis women), or that only a small subgroup of them are, is empirically wrong [2]. (Take the gender CV bias; also, trans-feminine people tend to downplay, not over-stress their discrimination because of internalized transphobia. [10]) Also, the ambivalent relation of ‘TERFs’ to vagueness (rejection while building it up as a danger), i.e. tolerant acceptability, is worth thorough philosophical investigation.

¹⁷ The analogue in the windows case would be asking *Did you water the plants?* when our actual interest is whether the flat will be flooded: It might make a ‘yes’ for the QUD more likely, but it doesn’t imply it.

¹⁸ A variation of [3, 4]’s probabilistic account might be a good place to look for this sort of connection.

¹⁹ I don’t claim my suggestion is the only one that makes sense of these cases. It just *also* makes sense of them, since it is more powerful.

Maybe this semantic-pragmatic interaction does not actually need introduction, but is a general property of natural language utterances that is sometimes blocked (e.g. by *all (windows)* as compared to *the (windows)* or more precise *person self-identifying as a woman* as compared to plain *woman*). [8]’s limited claims about non-maximal definite DPs could even be seen as a principled basis of communication at the semantics-pragmatics interface: The four truth/acceptability values can be reduced to two in several cases, which would equal a generalization to the most complex case. This would imply a certain position in the debate about semantic minimalism (the lexical meaning is held constant and may only be extended for pragmatic acceptability conditions via QUDs) and might be seen as specific a kind of pragmatic enrichment [14].

References

1. Burnett, H.: *Gradability in Natural Language: Logical and Grammatical Foundations*, Oxford Studies in Semantics and Pragmatics, vol. 7. Oxford University Press (2016)
2. Calderon-Cifuentes, P.A.: *Trans Discrimination in Europe. A TGEU analysis of the FRA LGBTI Survey 2019*. TGEU (2021)
3. Feinmann, D.: *Lexical Illusions, (Non-)Maximality, and Invisible Gaps*. Ph.D. thesis, University of Sheffield (2020)
4. Feinmann, D.: *Distinguishing Homogeneity From Vagueness*. *Journal of Semantics* **38**(4), 667–679 (12 2021)
5. Ginzburg, J.: *The Interactive Stance: Meaning for Conversation*. Oxford University Press UK (2012)
6. Haslanger, S.: *Gender and race: (what) are they? (what) do we want them to be?* *Noûs* **34**(1), 31–55 (2000)
7. Haslanger, S.: *Ideology, Generics, and Common Ground*. In: Witt, C. (ed.) *Feminist Metaphysics: Explorations in the Ontology of Sex, Gender and the Self*, pp. 179–208. Dordrecht (2011)
8. Haslinger, N.: *Non-maximality and vagueness: Revisiting the plural sorites paradox*. In: *Semantics and Linguistic Theory* 32. pp. 63–82 (2022)
9. Hyde, D., Raffman, D.: *Sorites Paradox*. In: Zalta, E.N. (ed.) *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2018 edn. (2018)
10. Jackman, K.B., Dolezal, C., Bockting, W.O.: *Generational differences in internalized transnegativity and psychological distress among feminine spectrum transgender people*. *LGBT Health* **5**(1), 54–60 (2018)
11. Jenkins, K.: *Amelioration and inclusion: Gender identity and the concept of woman*. *Ethics* **126**(2), 394–421 (2016)
12. Križ, M., Spector, B.: *Interpreting plural predication: Homogeneity and non-maximality*. *Linguistics and Philosophy* **44**(5), 1131–1178 (2021)
13. Malamud, S.A.: *The meaning of plural definites: A decision-theoretic approach*. *Semantics and Pragmatics* **5**(3), 1–58 (2012)
14. Recanati, F.: *Pragmatic enrichment*. In: Fara, D., Russell, G. (eds.) *Routledge Companion to Philosophy of Language*, pp. 67–78. Routledge (2010)
15. Roberts, C.: *Information structure: Towards an integrated formal theory of pragmatics*. *Semantics and Pragmatics* **5**, 1–69 (2012)