

Article

Not Strictly a Woman—QUD-Based Four-Valent Reasoning Discharges Lexical Meaning

Emil Eva Rosina ^{1,2,*}  and Franci Mangraviti ³

¹ Institute of Philosophy II, Ruhr University Bochum, 44801 Bochum, Germany

² Department of Czech Language, Faculty of Arts, Masaryk University, 602 00 Brno, Czech Republic

³ Department of Humanities, Social and Political Sciences, ETH Zürich, 8092 Zürich, Switzerland; legallyfranci@gmail.com

* Correspondence: emil.rosina@phil.muni.cz

Abstract

We 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’). We extend existing approaches to non-maximality 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 unacceptable cases of *x is a woman*, there are also (iii) merely pragmatically acceptable cases (‘like a woman with respect to the QUD’) as well as (iv) truly vague ones. We 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; polysemy; context dependency; vagueness



Academic Editors: Valentin Goranko and Wesley H. Holliday

Received: 28 June 2025

Revised: 5 November 2025

Accepted: 1 December 2025

Published: 11 December 2025

Citation: Rosina, E.E.; Mangraviti, F. Not Strictly a Woman—QUD-Based Four-Valent Reasoning Discharges Lexical Meaning. *Logics* **2025**, *3*, 16. <https://doi.org/10.3390/logics3040016>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

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. We 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

- (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 (identifies as a woman) wants to join. (x = Tammy) **invite!**
 - c. Non-binary person Noa (transmasculine, does not identify as a woman, mostly read female, suffering from imposter syndrome) wants to join. (x = Noa) **invite!**
 - d. Trans man Mario (does not identify as a woman, transitioned as a teenager, always read male) wants to join. (x = Mario) **invite?**
 - e. Assigned-male-at-birth Quinn (recently started questioning their gender, is unsure whether they identify as a woman, always read male and masculine-presenting) wants to join. (x = Quinn) **invite?**

- f. Cis man Chris (has never questioned his gender, does not identify as a woman, perceived as masculine and determinate) wants to join. ($x = \text{Chris}$) **do not invite**.
- g. $\llbracket x \text{ is a woman} \rrbracket = 1$ iff ... ?!

We 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 our illustration will keep the scenario constant, the solution cannot be ordinary polysemy resolution in dependence of the context either. We 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 person Noa's (1-c) positions in the Glass ceiling scenario. Section 2.1 discusses the political problem, Sections 2.2 and 2.3 provide the relevant linguistic tools, and Section 2.4 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 unacceptability. 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').

Section 3 focuses on the cases of trans man Mario and questioning Quinn as truly vague cases. We expand the analogy to [1] who distinguishes and brings together context-dependency that is resolved as soon as the QUD and other contextual parameters are fixed [2] from genuine vagueness where the QUD leaves open several different ways of partitioning the logical space [3]. We argue that this two-fold approach is straightforwardly applicable to Quinn and Mario and not restricted to the cases in plural semantics discussed by [1]. Again, the first subsection, Section 3.1, makes the political case, the second one, Section 3.2, provides the linguistic background, and the third subsection, Section 3.3, brings them together for a solution. Mario (1-d) and Quinn (1-e) are modeled 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 an apparent Sorites paradox (introduced in Sections 3.1 and 3.2) that might seem to lead to inviting cis man Chris as well. Together, Sections 2 and 3 show that a generalized version of [1]'s tools straightforwardly makes sense of optimal usage of gender terms and provide an argument against slippery slope worries about trans inclusion. 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 we model all of the intuitions in (1).

Finally, Section 4 addresses three loose ends: Section 4.1 shows that our account can also be used to model bad practice. Section 4.2 contains some pointers to a fruitful probabilistic extension that has the potential of answering worries concerning a strict implication from 'woman' to 'low on patriarchal hierarchies'. In Section 4.3, we conclude with possible implications of our account at the semantics–pragmatics interface in the case of other less politically loaded predicates.

1.3. Methodological Remark

This paper is methodologically non-standard both from a gender-philosophical as well as a linguistic point of view. We are applying a framework that has been developed for

non-maximal determiner-phrases (like *the windows*, as introduced in Section 2.2) to what seems to be a polysemy problem. Existing accounts of polysemy did not seem to be able to model how we *would like to* use gender terms like *woman*. If we had not considered that a problem, we would not have had any motivation to point to the parallel, because existing accounts of polysemy are able to deal with politically less loaded cases. That said, we think that the gender philosophical debate about context-dependency (see [4] a.o.) is based partly on intuitions that can be captured more precisely with existing linguistic tools, and we understand this paper as a step in this direction. This can, in turn, inform linguistics, since our claims do not have to be restricted to the politically relevant cases, see Section 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 do not make sense of the politically relevant cases. As for the original political motivation, we are confident that the account presented here provides the tools for better language policies, but we do not focus on how to put it into practice (e.g., how to actually phrase an invitation or how to decide in the vague cases in practice, but see note 13 and Section 3.3 for pointers). We also leave the relation to philosophical gender models to future gender-philosophical work.

2. Like a Woman—Truth and Issue-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’,⁴ *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’.⁵ Since, for many people, their gender self-identification does not change in dependence of the context,⁶ 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’ transmasculine people like Noa in (1-c), repeated below as (3-c), who was (let us assume) assigned female at birth and raised in a feminizing way.

- (3) a. **Scenario glass ceiling**
 b. Trans woman Tammy (identifies as a woman) wants to join. (x = Tammy) **invite!**
 c. Non-binary person Noa (transmasculine, does not identify as a woman, mostly read female, suffering from imposter syndrome) wants to join. (x = Noa) **invite!**

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 (3-b) is able to join the event (independent of her clothing, body features, etc.). At the same time, something seems wrong if Noa (3-c) is excluded even though they are subject to the kind of discrimination which the event targets: They are perceived by their coworkers in a way that makes them treat Noa like they treat women, i.e., (possibly just subconsciously)

not very well. Additionally, Noa was perceived as a woman and treated as intellectually inferior already in school, which—let us assume—has left traces in their self-esteem. Noa's treatment makes it such that they suffer from many of the same 'glass ceiling problems' as their women coworkers—regardless of how they identify.⁷

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. It is not enough to say that what counts as a woman in certain circumstances varies (e.g., $woman_3$ = 'low on the patriarchal hierarchy' in the Glass ceiling case (1-a) but $woman_1$ = 'self-identifies as a woman' amongst friends). 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.⁸ However, in Noa's case, this happens on the basis of calling them a woman—even if only in this context. Meanwhile, many other options $woman_n$ draw on irrelevant factors and just render more combinations of exclusion and misgendering harm, (4-c,d):

- (4)
- a. $\llbracket x \text{ is a woman}_1 \rrbracket = 1$ iff x self-identifies as a woman
→ Tammy invited, **Noa not invited**, no-one misgendered.
 - b. $\llbracket x \text{ is a woman}_2 \rrbracket = 1$ iff x was assigned female at birth
→ **Tammy not invited**, Noa invited, **both misgendered**.
 - c. $\llbracket x \text{ is a woman}_3 \rrbracket = 1$ iff x is low on the patriarchal hierarchy
→ Tammy and Noa invited, **Noa misgendered**.
 - d. $\llbracket x \text{ is a woman}_4 \rrbracket = 1$ iff x has an F in their passport
→ **invitation and misgendering depend on law and bureaucracy**.

Inviting everyone who is low on the patriarchal hierarchy ($woman_3$) picks out the right group for the Glass ceiling event in theory—the only harm done here (to Noa) is the association with the word *woman*. However, on the practical side, the question remains how we are going to get this group to attend the event—using which wording. Since *woman* seems to be a bad option at first glance (due to misgendering Noa), 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*). We reject this option for most events because—while avoiding problems with word meaning—it is open to attack with political arguments from practicability and classism.

First, not everyone who is subject to gender-based discrimination knows what is meant by this phrase. For example, someone who has not been in touch with feminist writing might identify as a woman and be subject to this kind of discrimination but not notice the uncommon phrase on a flyer or think 'some academic event that I wouldn't get anyway'. We believe that at least some events against gender-based discrimination should also (maybe specifically) target this group of people. Otherwise, we risk getting into a spiral of fewer and fewer emancipatory resources for a certain group of people: they will be more and more vulnerable to gender-based discrimination because they are not provided with options to name and see their oppression (a form of hermeneutical injustice [5]).

What is more, not everyone who is subject to gender-based discrimination knows that they are, in the sense that they would agree 'yes, I have systematic disadvantages because I am (or because I am perceived as) a woman'. Two groups of people come to mind who might not see it like this pre-attendance: (i) People who have never questioned their being assigned female at birth (mostly cis women) and for example think something like 'Men

and women are just different. Yes, I have my specific struggles in the household and with the kids, but my husband has his issues due to being a man as well. I wouldn't want to switch places with him. Overall, I think I'm privileged, getting to spend so much time with my kids. I like being emotional and dressing up.' (ii) Freshly cracked or questioning transfeminine people. Since they have been assigned male at birth, many of them have the impression of having been privileged based on their perceived and assigned gender up to this point. While these possible privileges commonly turn out to have been outweighed by the costs of being trans, and being transfeminine in particular, even when this identity was not conscious, this is a self-realization that is not a given—especially in an environment of transmisogyny where transfeminine people are portrayed as threats and imposters.

We would like our invitation policies to be such that they do not assume that every individual is able to optimally distance themselves from internalization of toxic narratives. Hence, the invitation has to make clear that self-identification as a woman or gender minority is a no-brainer with respect to being invited. While an event like Glass ceiling should optimally provide attendants with the means to critically reflect on all of the above points (while taking their initial perspective seriously), it cannot assume that participants have already undergone such reflection pre-attendance. Therefore, a phrasing like 'this event is for everyone who is subject to gender-based discrimination or gender bias' is a form of gatekeeping in many contexts. If we are not understood and no-one joins our event, there is no use in a theoretically good choice of words.

Our alternative suggestion introduced in the rest of the paper does involve the term *woman* in the invitation (while also solving the misgendering problem for Noa). This involves a commitment that the event is relevant to women (and others). Forefronting self-identification in this way is a means to signal the above two groups of people that the organizers consider the issues discussed at the event in question relevant to them. This can be a more welcoming environment than one that leaves the interpretation of who is seen as 'subject to gender-based discrimination' up to imagination.⁹

If we want to stick to using *woman* in some contexts, then, modeling 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., [6])—seems to linguistically involve 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. The overall challenge is to resolve a deadlock between three bad options discussed so far: misgendering Noa by inviting them calling them a woman, not inviting Noa at all and keeping deserved resources from them, or using inaccessible phrasing and risking not being understood. We 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: Non-Maximality and Vagueness, Haslinger (2022)

This paper exploits a structural parallel between the (language-)political problem introduced in Sections 1.1 and 2.1 and a non-political linguistic phenomenon—the non-maximality of definite determiners like English *the*. Therefore, we sketch this phenomenon quickly here and informally introduce [1]'s treatment of non-maximality which we are going to transfer to the distinct phenomenon of gender predicates. Consider (5):

- (5) The windows are open.

Assume that a given salient flat has 10 windows. Intuitions about how many of these have to be open for (5) to be true vary between individuals and are unclear even for many individual speakers. Unlike *all windows are open*, (5) does not require 10 of 10 windows to be open, at least not in all circumstances. Two previous suggestions for a semantics/pragmatics of *the* in these cases include issue-based approaches like [2] and vagueness-based approaches like [3].

The rough idea of issue-based approaches to non-maximality is that the unclear intuitions on sentences like (5) are due to underspecification of contextual parameters. First, it is impossible to judge the truth of (5) without knowing what the Question under Discussion ('QUD')—simplifying, the question which is to be addressed by the current move in the conversation—is. (We will say more about QUDs in the following Section 2.3) Second, even if this goal is fixed, the issue parameter, fixing a division of the logical space which is determined by the context, has to determine which worlds 'cluster together' with respect to the task of resolving the QUD. For the purpose of fixing both as much as needed for the case at hand, let us assume the scenario in (6-a).

- (6) a. **Scenario thunderstorm:** Ann and Bea left their flat. A thunderstorm is coming. A total of 4 of the 10 windows face the direction of the wind and are 2 m wide; the other 6 are small bathroom windows. Unfortunately, Ann left exactly the four most crucial windows open before leaving. Ann: *Oh no!...*
- b. *The windows are open!*¹⁰ **QUD:** Will our flat be flooded?
Issue parameter: The 4 large windows being open has the same effect on the flat during a thunderstorm as 10 windows being open. It is unclear whether this is also the case for three, two, or one large window being open. If no large windows are open, the flat will be fine. Bathroom windows do not matter.

The idea of issue-based accounts of non-maximality (like [2]) is that non-maximal readings arise whenever the issue parameter is such that some other configuration (of open and closed windows in our case) has the same effect for resolving the QUD as the maximal case (where all windows are open). A strong argument in favor of such a treatment is the fact that *which* window is open—a big one or a small bathroom window—can make a difference, not just their number. Even if six bathroom windows (and no others) were open, (5) would end up unacceptable (our terminology; [1] speaks of 'pragmatic truth/falsity'). Once both the QUD and issue parameter are specified enough to provide a clear clustering or non-clustering of the given case (4 large windows open) with the maximal case (10 windows open), the sentence is either clearly acceptable or clearly unacceptable.

The rough idea of vagueness-based approaches to non-maximality (like [3]) is to explain the unclear intuitions on (5) instead as the same kind of phenomenon we encounter when judging whether someone is tall or not: There is a whole range of tolerant ('border-line') cases of people who are neither tall nor not tall, and this does not have to be an issue of underspecification. Similarly, there could be a range of situations with open windows where it is simply not clear (and not resolvable) whether the given set of open windows count as *the* windows already or not.

Ref. [1] suggests that these two classes of accounts—issue-based and vagueness-based—are not, in fact, exclusive. Instead, there are cases of non-maximality that can be attributed to the issue parameter only (in a resolvable way) and cases where vagueness is needed on top to explain intuitions. The distinction between vague and non-vague non-maximal contexts is modeled in terms of the relation of the issue parameter to the QUD. If the issue parameter in (6) was perfectly specified, the answer to the flooding QUD would be clear for every possible configuration of open and closed windows. This is unrealistic for (6) and most other natural conversations—the context leaves open what happens in the

cases of one, two, or three large windows being open. In this sense, (6) is a context that can, in principle, give rise to vagueness. However, this does not matter for the case where four large windows are open, as described, because it does not fall into the tolerance range. The unclear intuitions on (5) when 4 windows are open can be entirely attributed to the issue parameter which—when specified as in (6-b)—clearly clusters together the 4/10 and the 10/10 case. We will later suggest that this is parallel to Noa’s situation in the Glass ceiling scenario. Consequently, we suggest that there is an overarching phenomenon of ‘issue-based acceptability’ which goes beyond non-maximality.

Finally, ref. [1] suggests that for vagueness to arise, there must be several potential values for the issue parameter, which stand in a non-transitive similarity relation. (We add that this must also be relevant for the case at hand.) For example, it is not realistic that only exactly four open large windows lead to flooding, but it may well be unclear to everyone in a given conversation—as fixed in (6-b)—what happens in the case of one, two, or three open large windows. The context simply does not provide the means to determine the exact issue parameter, as often occurs in natural conversation. Hence, the issue parameter in (6-b) is not maximally specified. Therefore, to model non-maximality in a case that is otherwise like (6) but where, say, two large windows are open, we need tolerance on top. In this case, (5) is *vaguely* acceptable. We will later suggest that this is analogous to Quinn’s and Mario’s situation in the Glass ceiling scenario.

To sum up, ref. [1] distinguishes context-dependency that is resolved as soon as the issue parameter is fixed (enough) from genuine vagueness and combines these two factors into a model. Vagueness arises when the QUD leaves open several different ways of partitioning the logical space, and thus several different issue parameters (and the given case falls into the tolerance range—otherwise, the potential vagueness stays unrealized). In this paper, we transfer this idea to the Glass ceiling case even if it does not involve non-maximality.

2.3. Linguistic Background: Issue-Based Acceptability

In the remainder of Section 2, we motivate the suggested parallel between issue-based acceptability with respect to non-maximal meanings of sentences like (6-b), repeated as (7-b), and cases like Scenario glass ceiling. After introducing issue-based non-maximality a bit more in this subsection, we will show that [1]’s solution to these cases can be transferred to Noa’s situation in the Glass ceiling case in Section 2.4.

While definite DPs like *the windows* in (7-b) give rise to the same (semantic) truth-conditions as universal quantifier phrases (*all windows*) in sentences like (7-b), the former are often accepted in non-maximal scenarios like thunderstorm.

- (7) a. **Scenario thunderstorm:** Ann and Bea left their flat. A thunderstorm is coming. A total of 4 of the 10 windows face the direction of the wind and are 2 m wide, while the other 6 are small bathroom windows. Unfortunately, Ann left exactly the four most crucial windows open before leaving. Ann: *Oh no!...*
- b. *The windows are open!* **QUD:** Will our flat be flooded?
Issue parameter: The 4 large windows being open has the same effect on the flat during a thunderstorm as 10 being windows open. It is unclear whether this is also the case for three, two, or one large window being open. If no large windows are open, the flat will be fine. Bathroom windows do not matter.

The above has been explained via what we will be calling (pragmatic) acceptability conditions, which can be much weaker than the (semantic) truth-conditions (‘literal meaning’) of a sentence, depending on features of the context. For example, refs. [2,8], building on [9], resort to the concept of Question under Discussion (‘QUD’), complemented here by [1]’s notion of the issue parameter, going back to [10], which we briefly introduced

in Section 2.2. Roughly, a QUD is the motivation for the current move in a conversation. For the case of Scenario thunderstorm, the salient QUD is the underlying question whose answer requires speakers to know the status of the windows—something like ‘Will our flat be flooded?’.¹¹ Given the context that we fix in (7), the answer is 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. The partitioning of the logical space relative to the QUD is done by an *issue parameter* which is sensitive to details of the context like the size of the windows in (7-a). While the QUD in this case is a polar question, so Ann and Bea only have to find out whether they are in a yes- or a no-world, the issue parameter is much more complex. If optimally specified, it would provide a complete partition (e.g., ‘exactly 4 large windows being open clusters with 10 windows being open; 3 or less large windows being open clusters with no windows being open; the status of small windows has no influence whatsoever’). In practice, a fully specified issue parameter is rare—but the only thing that matters for cases like (7) is that it is specified enough to determine the given case.

Once enough of this information is known to everyone, it does not matter anymore to the current purpose whether the sentence ‘All windows are open’ or ‘Four windows are open’ was used. Put simply, 4/10 windows is *like* 10/10 windows for the flooding QUD, given the information on the issue parameter in (7). (7-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 dependency on the QUD and issue parameter is a plausible root of the acceptability of (7-b) in Scenario thunderstorm becomes clear in contrast with (8), where the QUD is different, and (9) where only the issue parameter is different:

- (8) a. **Scenario paint:** Ann and Bea left their newly painted flat. The paint requires maximal ventilation. Four of ten windows are open. Bea: *Did you leave everything like the instructions say?* Ann: *Yes! ...*
 b. ??*The windows are open!* **QUD:** Will the paint dry?
Issue parameter: At least 9 of 10 windows must be open for the paint to dry sufficiently.

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 according to the issue parameter; i.e., these cases are *not* in the same partition with respect to this QUD. But even if the QUD stays the same, the issue parameter can make a big difference:

- (9) a. **Scenario thunderstorm II:** Ann and Bea left their flat. A thunderstorm is coming. Luckily, none of the 10 windows face the direction of the wind and all of them are protected from an outside roof. Only if all of them are open is any damage to the flat realistic. Ann left 4 of these windows open before leaving. Ann: *Oh no!...*
 b. #*The windows are open!* **QUD:** Will our flat be flooded?
Issue parameter: Four windows open has the same effect on the flat during a thunderstorm as no windows being open. Only 10 windows being open causes flooding.

Scenarios (9) and (8) together show that the relevant factor for non-vague non-maximal readings is an interaction between QUD and the issue parameter. For the purpose of simplicity, we will from now on talk of ‘issue-based’ or ‘issue-dependent acceptability’ when referring to this interplay.

2.4. Application: Issue-Dependence in the Glass Ceiling Case

Extending issue-based acceptability conditions to cases with gender predicates solves language policy problems. We suggest that the more general equivalent of a non-maximal scenario is one that does not make a sentence semantically true but pragmatically acceptable, because of QUD answer equivalence to the semantical-truth-evoking case.¹² Let us reconsider Scenario glass ceiling in this light:

- (10) a. **Scenario glass ceiling** **QUD:** Is x low on the patriarchal hierarchy?
Issue parameter: Women are low on the patriarchal hierarchy. Anyone who is routinely affected by disadvantages due to gender CV bias, being mansplained at, gendered credibility deficit, gendered imposter syndrome, sexist remarks, etc., in the workplace is also low on the patriarchal hierarchy. Anyone who only occasionally experiences some of these things is in the tolerance range.
- b. Trans woman Tammy (identifies as a woman) wants to join. (x = Tammy) **invite!**
- c. Non-binary person Noa (transmasculine, does not identify as a woman, mostly read female, suffering from imposter syndrome) wants to join. (x = Noa) **invite!**
- d. x is a woman. **True for Tammy (10-b), false but acceptable for Noa (10-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 (10-d) is semantically true in subscenario (10-b). Adding issue-based acceptability allows Noa in as well without (clearly) misgendering them and makes (10-d) acceptable in subscenario (10-c). Noa is *like* a woman with respect to this specific issue, but unlike one in many other respects. Regardless of general similarities and differences between Tammy and Noa, the situation is such that they are both gender-discriminated against and, hence, in the same partition with respect to the QUD.

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 in this case ('What is Noa's gender?'). For this new QUD, it does make a difference whether Noa is a woman or non-binary. 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.¹³

The central move here, in relation to the dilemma sketched in Section 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 patriarchal hierarchy' and serves as a QUD now that it is only indirectly related to gender (in that every or practically—see Section 4.2—every $woman_1$ is low on the patriarchal hierarchy in the actual world).

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 we have already established $woman_1$ = ‘self-ID woman’ as the semantic truth criterion, inviting trans women like Tammy to the Glass ceiling event is not in question at this point, even if she does not take hormones and is butch-presenting.¹⁴ There are, however, cases where even trans-inclusive intuitions about whether to invite someone to the Glass ceiling event are unclear, and we do not want that to lead to inviting everyone in.¹⁵ In Sections 3.2 and 3.3, we are going to show that the tools from literature on vagueness suffice to stop these worries. First, we consider the political case.

- (11) a. **Scenario glass ceiling** **QUD:** Is x low on the patriarchal hierarchy?
Issue parameter: Women are low on the patriarchal hierarchy. Anyone who is routinely affected by disadvantages due to gender CV bias, being mansplained at, gendered credibility deficit, gendered imposter syndrome, sexist remarks, etc., in the workplace is also low on the patriarchal hierarchy. Anyone who only occasionally experiences some of these things is in the tolerance range.
- b. Trans man Mario (does not identify as a woman, transitioned as a teenager, always read male) wants to join. (x = Mario) **invite?**
- c. Assigned-male-at-birth Quinn (recently started questioning their gender, is unsure whether they identify as a woman, always read male and masculine-presenting) wants to join. (x = Quinn) **invite?**
- d. Cis man Chris (has never questioned his gender, does not identify as a woman, perceived as masculine and determinate) wants to join. (x = Chris) **do not 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, etc.) 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 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, this leads to the worry that 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.

Note that both ‘being subject to gender-based discrimination’ and ‘being low on the patriarchal hierarchy’ are arguably *scalar* properties: At one end of the scale, Tammy and Noa are highly discriminated against, which allowed us to treat patriarchal discrimination like a clear yes or no question up to this point. This treatment is now threatened by the existence of cases like Mario and Quinn. Chris is designed to be a clear case of very low gender discrimination.¹⁶

3.2. Linguistic Background: Vagueness

Things get more complicated than the three-fold distinction ‘semantic truth–mere pragmatic acceptability–unacceptability’ once we look at cases like (12), a variation of Scenario paint with 70/100 open windows:

- (12) a. **Scenario paint 2:** Ann and Bea left their newly painted luxury estate. The paint requires maximal ventilation in order to dry. A total of 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?

Issue parameter: A total of 95–100 open windows clearly suffice for the paint to dry; 1–50 windows clearly do not suffice; for 51–94 windows, it is unclear.

While semantic truth requires all salient windows in the given situation to be open, and no open windows makes (12-b) clearly unacceptable, 99/100 windows constitutes an acceptable case here relative to the QUD ‘Will the paint dry?’ and the given limitations on the issue parameter. However, intuitions on (12-b) in Scenario paint 2 with 70/100 open windows might vary and lead to unclear judgments. This is strongly reminiscent of standard cases of vagueness, like with *x is tall*, that give rise to a version of an apparent Sorites paradox (see, e.g., [14]): someone might be tempted to accept both premises (13-a,b) without accepting the conclusion (13-d) which classically follows (given repeated application of (13-b) to a bunch of people).

- (13)
- a. Susi, who is 2 m tall, is tall.
 - b. Any person who is only a millimeter less tall than a tall person is also tall.
 - c. Mimi, who is 1.999 m tall, is tall. (etc.)
 - d. Max, who is 1.4 m tall, is tall.

Sorites reasoning can also be applied to opening one window after the other: A particular speaker judges (12-b) unacceptable 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.

Of course, research in semantics has long provided solutions to this apparent paradox, e.g., in the form of vagueness-based accounts to plurals like [3]. (We are going to present the solution as integrated with issue-based accounts like [1] does in order to ultimately deal with the complete Glass ceiling case, but most accounts of vagueness could do the job when it comes to the apparent Sorites.) Treating vague cases like (12-b) in Scenario paint 2 as tolerantly acceptable (the forth ‘truth/acceptability value’ we are going to use)¹⁷ stops the modus ponens that enables the apparent Sorites paradox: a single window out of 100 does not make a difference between strict acceptability and unacceptability, but it might make a difference between strict acceptability and tolerant acceptability, or between tolerant acceptability and unacceptability. Following [1], building on [3], tolerant acceptability arises when the combination of the QUD and (what we know about) the issue parameter involves vagueness, i.e., when the explicit context justifies several different possible values of the issue parameter. In the case at hand, this means that there are several legitimate interpretations of what ‘counts as’ enough open windows/dry-enough paint when the tolerance range of 51–94 open windows is concerned, and consequently, of which cases are in the same partition as the strictly true/acceptable cases. What we 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 making overt one clear maximally specified issue parameter).

The tolerant cases are the ones that ‘are like’ the strictly true/acceptable ones with respect to some specification of the issue parameter, but not with respect to another one. Applied to (12-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 will not dry sufficiently with 70/100 windows open (and so this case partitions with the clearly unacceptable cases, like 0/100 or 2/100 windows open, and is not even pragmatically acceptable). The way [1] 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 and issue parameter are fixed. (However, refs. [2,3,8] present their accounts as full solutions to non-maximality. Ref. [1] points to their combinability.)

3.3. Application: Tolerant Acceptability in the Glass Ceiling Case

Ref. [1]’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 unacceptable 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.

- (14) a. **Scenario glass ceiling** **QUD:** Is x low on the patriarchal hierarchy?
Issue parameter: Women are low on the patriarchal hierarchy. Anyone who is routinely affected by disadvantages due to gender CV bias, being mansplained at, gendered credibility deficit, gendered imposter syndrome, sexist remarks, etc., in the workplace is also low on the patriarchal hierarchy. Anyone who only occasionally experiences some of these things is in the tolerance range.
- b. Trans man Mario (does not identify as a woman, transitioned as a teenager, always read male) wants to join. (x = Mario) **invite?**
- c. Assigned-male-at-birth Quinn (recently started questioning their gender, is unsure whether they identify as a woman, always read male and masculine-presenting) wants to join. (x = Quinn) **invite?**
- d. Cis man Chris (has never questioned his gender, does not identify as a woman, perceived as masculine and determinate) wants to join. (x = Chris) **do not invite.**
- e. x is a woman. **Tolerantly acceptable for Mario (14-b) and Quinn (14-c), unacceptable for Chris (14-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 underspecification) and, in practice, to let Mario and Quinn decide for themselves in most cases. The resistance to giving a clear positive or negative judgment arises exactly in those cases where the underlying issue is too vague itself to even judge pragmatic acceptability. In other words, the context gives rise to several possible values of the issue parameter, in- and excluding different parts of the tolerance range, just like in the 70/100 windows case. We suggest that [1]’s tools can therefore be transferred to model faultless political disagreement.

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: 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 sex at birth 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. For example, they might feel alienated from male-gendered academic networks or feel addressed by sexist remarks. 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: Mario's self-socialization as a man could have 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 at a theoretical 'gender is a construct' level.

It is 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 questioning is for them. That would be analogous to 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 (14-d) to the Glass ceiling event. Given a basic understanding of vagueness, this apparent Sorites paradox never gets off the ground, even in the presence of borderline cases.

Note that adding scalarity to the picture in Section 3 (in order to deal with vagueness) does not by itself annihilate the need for (resolvable) issue-dependence of the 'being like a woman in some respect' sort as described in Section 2: a scalarity-based solution alone, without separation of truth and acceptability, remains one-dimensional in the sense that 'the more you are gender-discriminated, the more you are invited, the more you are a woman', 'the more you are a woman, the more you are gender-discriminated, the more you are invited', etc. While this is helpful for people who partly identify as a woman or who identify as partly-a-woman (and are also partly subject to gender-based discrimination), it is not helpful for people who do not identify as a woman at all (like Noa) but who are to a high degree subject to gender-based discrimination (like Noa). A uni-dimensional scale is also only helpful (for some people) in the Glass ceiling scenario in isolation—there are, for example, also medical contexts, which draw on completely different properties. A single scale cannot make sense of the way in which the context matters to determine which properties are relevant. We argue that both issue-dependency in the sense of clear equality with respect to the QUD ('like a woman with respect to discrimination?') and vagueness are needed to cover all cases in a satisfactory way: For Noa, we have to be able to explain why they are clearly invited even though it is clearly false that they are a woman—this can not be done with vagueness alone. For Quinn and Mario, we need to be able to explain their borderline status—which cannot be done with issue-dependency alone. (Ref. [1] argues the same for the case of non-maximal determiners.)

Separating these two pragmatic mechanisms has helped us deal with the multi-dimensionality of gendered life without treating all dimensions on par: differentiating between sex, gender, and self-identification as several equal dimensions of a single lexical entry, *woman* (or other gender terms), is seen as problematic in current gender philosophy. For example, ref. [15] critiques naive contextualism—the idea that someone is a woman in a medical context but a man in a colloquial context, etc.—and [16] discusses problems with family resemblance accounts of gender—the idea that one is more or less of a woman depending on how many stereotypical properties one has—which weigh all dimensions. All properties that are practically associated with gender can still be discussed properly under our account: gender, equaling self-identification, as truth-conditional, and all others as different combinations of QUD and issue parameter. Hence, we account for the scalar nature of some of these properties while also separating them more clearly in a way that is sensitive to the practical needs of a situation. Before we move on to some loose ends, Table 1 sums up our solution to Scenario glass ceiling.

Table 1. Truth/acceptability conditions of *x is a woman* in Scenario glass ceiling.

	iff...	Tammy	Noa	Mario	Quinn	Chris
true	<i>x</i> self-identifies as a woman	✓	×	×	?	×
strictly acceptable	<i>x</i> is in the same partition with respect to the QUD as <i>y</i> , and <i>y</i> self-identifies as a woman ¹	✓	✓	?	?	×
tolerantly acceptable	<i>x</i> is in the women-partition with respect to some interpretation of the Issue parameter but not with respect to some other	×	×	✓	✓	×
unacceptable	neither of the above	×	×	×	×	✓
invited?		✓	✓	?	?	×

¹ 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.

4. Open Issues

4.1. Modelling Ethically Bad Use of Gender Terms

Our four-way issue-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 do not consider trans women ‘real women’, i.e., 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 our four-valent framework is that it captures many TERFs’ 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 combinations of QUDs and issue parameters are adequate, and the pragmatic extension of *woman* must therefore vary.¹⁸

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 is a woman*₁ = ‘self-ID woman’, because then being a woman would not imply being affected by the reason for the event.¹⁹ While we take discrimination of women (trans and cis) to be an established fact at the group level, the way we have put this so far may make it seem like a single woman₁ (trans or cis) who is not de facto discriminated against at all at work could destroy the whole line of argumentation. While we are not aware of such a case, this would not constitute a problem for our account regardless, because a certain idealization (and, consequently, no strict implication in practice) is 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. For example, 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.²⁰

While this probability is extremely high in the Glass ceiling case, a structurally similar case with giving birth as the relevant QUD might not actually justify the choice of QUD–woman₁ combination: While almost all women₁ are susceptible to the glass ceiling, being a woman₁ is only a weak predictor for giving birth, and even for being able to give birth (not only considering trans people, as many cis woman cannot give birth—e.g., after their menopause). While there is at least a plausible idealization of a unidirectional implication from strict semantic truth to a positive answer to the QUD in the Glass ceiling case, the Giving birth case might be less clear because of this weaker causality. Hence, maybe we should actually not use *women* on the invite in this case.

4.3. Consequences for the Pragmatics–Semantics Interface

We 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.²¹ Nor does the syntactic category intuitively play a role when we 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.

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 the more precise *person self-identifying as a woman* as compared to the plain *woman*). Ref. [1]’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 a specific kind of pragmatic enrichment [22].

Author Contributions: Conceptualization, E.E.R. and F.M.; methodology, E.E.R. and F.M.; writing—original draft preparation, E.E.R.; writing—review and editing, E.E.R. and F.M. All authors have read and agreed to the published version of the manuscript.

Funding: The first author’s research was funded by the German Research Foundation, DFG, grant number 397530566. The second author’s research was funded by the German Research Foundation, DFG—534806587.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Acknowledgments: We thank Jordan Chark, Peter Sutton, Nina Haslinger, Kristina Liefke, Elin McCready, the audience of an Internal Workshop of the Vienna Forum for Analytic Philosophy as well as three anonymous ESSLII reviewers and two anonymous MDPI reviewers for their helpful comments on different stages of this work.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

QUD	Question under Discussion
TERF	Trans-exclusionary ‘radical feminist’

Notes

- 1 We are not making any assumptions about the nature of this choice/‘choice’.
- 2 We take misgendering to be equivalent with not respecting someone’s gender self-ID.
- 3 We do not claim that our 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. We 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, we point to a set of possible tools for modeling (what is wrong with) non-optimal actual uses of gender terminology, see Section 4.1.
- 4 By ‘low on the patriarchal hierarchy’ we mean ‘low on the gender-related hierarchy that comes with patriarchal societies’, or roughly ‘highly vulnerable to discrimination because of (actual or perceived) gender’. We use ‘subject to patriarchal discrimination’ interchangeably.
- 5 We are 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, etc.) often associated with gender. However, neither of them is a reliable predictor of self-ID. We are not making any claims about the metaphysics of gender here.
- 6 We are aware that gender-fluidity contrasted with retrospective self-ID complicates things. Nonetheless, we have to exclude the time component for matters of space.
- 7 A variant of Noa where they are non-binary and transfeminine (so assigned male at birth) would lead to the same outcome, stressing current, external discrimination more and previously internalized discrimination less.
- 8 Yes, trans women are low on the patriarchal hierarchy. See also Sections 4.1 and 4.2 for the importance of this implication.
- 9 Even if none of the internal narratives discussed here play a role in a given context of a Glass ceiling meeting, announcing most feminist events with the less accessible phrasing gives rise, in a toxic political society, to narrations along the lines of ‘wanting to cancel the existence of women’ and ‘queerfeminists building their elitist bubble’. While the latter is a misconception of most queerfeminist environments, a true core of it (that can then be exploited for further hate and discrimination) can be the following: using only the less accessible phrasing merely *declares* contexts to be such that it should be common ground what ‘subject to gender-based discrimination or gender bias’ means. We argue that, rather than declaring contexts such, contexts should be *made* such by including and empowering people who do not share a feminist common ground on their oppression yet.
- Of course, if we know it to be very likely that none of the narratives above play a role in a given context, it is legitimate to phrase the invitation like this—for example, if we know that everyone in the target group at a given institute has already attended a first feminist event earlier that has prepared them to understand that they are meant by ‘people who are subject to gender-based discrimination or gender bias’. The problem we discuss in this paper, however, remains for all other cases.
- 10 Empirical support for non-maximal readings of the structurally analogous sentence ‘The squares are green’ comes from [7].
- 11 See e.g., [11] for a more complete picture of QUDs.
- 12 While [2] suggest that QUD-sensitivity has to be introduced by elements like *the* and is otherwise absent, there is independent discussion [12] of QUD-sensitivity as the default. QUD-*insensitivity* would have to be introduced, e.g., by universal quantifiers who would act as filters of QUDs. While the cases we discuss do obviously not involve definite determiners, we suspect that the parallel is not coincidental and should be traced back to general principles of context- and QUD-sensitivity.
- 13 We thank an anonymous reviewer for pushing us to comment on this point with regard to Quinn and Mario (Section 3.3). We take this issue to be even more pressing in the strictly pragmatically acceptable case of Noa. Being a trans-masculine non-binary person themselves, the first author of this paper would not feel misgendered by the sort of pragmatic acceptability involved here, if they were Noa. On the contrary, they 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, we 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.
- Practically speaking, our account is probably best realized by making the pragmatically added group explicit in order to avoid ambiguities in a world where different people use very different notions of *woman*—for example, by inviting ‘women and everyone else who is discriminated against based on gender’. This is, however, not the main focus of our paper.

- 14 Instead, still excluding Tammy from the event would come down to not granting her first person authority based on what she
says is her gender (cf. [13]).
- 15 ...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.
- 16 We treat the scale of ‘being low on the patriarchal hierarchy’ as linear for purely technical reasons in our model—it suffices to
provide the intuitively correct results for our cases. This does not involve any ontological or practical commitment that every
pair of individuals can be ordered in terms of gender-based discrimination. There may well be good reasons (for example
considerations from the intersections of gender with racialization, class, disability etc.) to refrain from such commitment. In
future work, this linear scale should itself be captured as multi-dimensional and complex.
- 17 We are not making any ontological or logical claims, this is just intuitive terminology.
- 18 This is by far not the only fallacy in TERF reasoning, of course: Their choice of lexical meaning does misgendering harm. The
belief that trans women are not suffering from patriarchy (enough/as much as cis women), or that only a small subgroup of
them are, is empirically wrong [17]. (Take the gender CV bias; also, trans-feminine people tend to downplay, not over-stress their
discrimination because of internalized transphobia [18].) 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.
- 19 The analog 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 does not imply it.
- 20 A variation of [19,20]’s probabilistic account might be a good place to look for this sort of connection.
- 21 We do not claim our suggestion is the only one that makes sense of these cases. It just *also* makes sense of them, since it is
more powerful. The only other linguistic account that we know of for applying QUD-dependence to predicates in analogy to
non-maximality—[21]—works substantially differently from ours and does not lead to the desired result for *women*.

References

- Haslinger, N. Non-maximality and vagueness: Revisiting the plural Sorites paradox. In Proceedings of the Semantics and Linguistic Theory 32, Mexico City, Mexico, 8–10 June 2022; pp. 63–82.
- Križ, M.; Spector, B. Interpreting Plural Predication: Homogeneity and Non-Maximality. *Linguist. Philos.* **2021**, *44*, 1131–1178. [[CrossRef](#)] [[PubMed](#)]
- Burnett, H. *Gradability in Natural Language: Logical and Grammatical Foundations*; Oxford Studies in Semantics and Pragmatics; Oxford University Press: Oxford, UK, 2016; Volume 7, p. 208.
- Cull, M.J. Queer Identities in Contexts. In *The Routledge Handbook of LGBTQ Identity in Organizations and Society*; Routledge: London, UK, 2024; pp. 365–376.
- Fricker, M. *Epistemic Injustice: Power and the Ethics of Knowing*; Oxford University Press: Oxford, UK, 2007.
- Jenkins, K. Amelioration and Inclusion: Gender Identity and the Concept of Woman. *Ethics* **2016**, *126*, 394–421. [[CrossRef](#)]
- Križ, M.; Chemla, E. Two methods to find truth-value gaps and their application to the projection problem of homogeneity. *Nat. Lang. Semant.* **2015**, *23*, 205–248. [[CrossRef](#)]
- Malamud, S.A. The meaning of plural definites: A decision-theoretic approach. *Semant. Pragmat.* **2012**, *5*, 1–58. [[CrossRef](#)]
- Roberts, C. Information Structure: Towards an integrated formal theory of pragmatics. *Semant. Pragmat.* **2012**, *5*, 1–69. [[CrossRef](#)]
- Groenendijk, J.A.G.; Stokhof, M.J.B. Studies on the Semantics of Questions and the Pragmatics of Answers. Ph.D. Thesis, University of Amsterdam, Amsterdam, The Netherlands, 1984.
- Ginzburg, J. *The Interactive Stance: Meaning for Conversation*; Oxford University Press UK: Oxford, UK, 2012.
- Haslinger, N.; Rosina, E.E.; Schmitt, V.; Wurm, V. On the relation between distributivity and maximality. *Semant. Pragmat.* **2025**, *18*, 1–23. [[CrossRef](#)]
- Bettcher, T.M. Trans Identities and First-Person Authority. In *You’ve Changed: Sex Reassignment and Personal Identity*; Shrager, L.J., Ed.; Oxford University Press: Oxford, UK, 2009.
- Hyde, D.; Raffman, D. Sorites Paradox. In *The Stanford Encyclopedia of Philosophy*, Summer 2018th ed.; Zalta, E.N., Ed.; Metaphysics Research Lab, Stanford University: Stanford, CA, USA, 2018.
- Saul, J. Politically significant terms and philosophy of language: Methodological issues. In *Out from the Shadows: Analytical Feminist Contributions to Traditional Philosophy*; Oxford University Press: Oxford, UK, 2012; pp. 195–216.
- Cull, M. Engineering Genders: Pluralism, Trans Identities, and Feminist Philosophy. Ph.D. Thesis, University of Sheffield, Sheffield, UK, 2020.
- Calderon-Cifuentes, P.A. *Trans Discrimination in Europe. A TGEU Analysis of the FRA LGBTI Survey 2019*; TGEU: Singapore, 2021.
- Jackman, K.B.; Dolezal, C.; Bockting, W.O. Generational Differences in Internalized Transnegativity and Psychological Distress Among Feminine Spectrum Transgender People. *LGBT Health* **2018**, *5*, 54–60. [[CrossRef](#)] [[PubMed](#)]
- Feinmann, D. Lexical Illusions, (Non-)Maximality, and Invisible Gaps. Ph.D. Thesis, University of Sheffield, Sheffield, UK, 2020.
- Feinmann, D. Distinguishing Homogeneity From Vagueness. *J. Semant.* **2021**, *38*, 667–679. [[CrossRef](#)]

21. Paillé, M. Truth-value gaps with integrative predicates. In Proceedings of the 41st West Coast Conference on Formal Linguistics. Cascadilla Proceedings Project, Santa Cruz, CA, USA, 5–7 May 2023; pp. 466–475.
22. Recanati, F. Pragmatic Enrichment. In *Routledge Companion to Philosophy of Language*; Fara, D., Russell, G., Eds.; Routledge: London, UK, 2010; pp. 67–78.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.