

ARTICLE

How to Solve the Gender Inclusion Problem

Cameron Domenico Kirk-Giannini

Department of Philosophy, Rutgers University–Newark
Email: camerondomenico.kirkgiannini@gmail.com

(Received 21 March 2024; revised 18 July 2024; accepted 30 September 2024)

Abstract

The *inclusion problem* for theories of gender arises when those theories inappropriately fail to include certain individuals in the gender categories to which they ought to belong. The inclusion problem affects both of the most influential traditions in feminist theorizing about gender: social-position accounts and identity accounts. I argue that the inclusion problem can be solved by adopting a *structured* theory of gender, which incorporates aspects of both social-position accounts and identity accounts. According to the theory I favor, an individual's gender is determined by their gender identity if they have one; otherwise, it is determined by their social position. My structured theory of gender offers a more direct solution to the gender inclusion problem than alternatives recently advocated by Barnes (2020), Jenkins (2023), and others. It also points the way to a simple solution to inclusion problems that arise at the level of gendered language.

Feminist accounts of gender hold that it is a social phenomenon conceptually and modally distinct from the biological phenomenon of sex.¹ Particular feminist theories of the social construction of gender generally fall into one of two classes. To a first approximation, *social-position accounts* suggest that an individual's gender is determined by the way they are thought about and treated by other members of society, while *identity accounts* suggest that an individual's gender is determined by their own internal sense of themselves as a woman, a man, or neither.

Despite their centrality in the feminist tradition of reflection on the nature of gender, both social-position accounts and identity accounts have recently been criticized as exclusionary. On the one hand, social-position accounts seem to rule out the possibility that someone who is thought about and treated as a man might in fact be a woman. This consequence has been thought exclusionary because it fails to classify certain trans women as women. On the other hand, identity accounts seem to rule out the possibility that someone who lacks a gender identity might have a gender. This consequence has been thought exclusionary because it fails to classify certain intellectually disabled

© The Author(s), 2025. Published by Cambridge University Press on behalf of Hypatia, a Nonprofit Corporation. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial licence (<https://creativecommons.org/licenses/by-nc/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original article is properly cited. The written permission of Cambridge University Press must be obtained prior to any commercial use.

women as women.² But if understanding gender as imposed on an individual externally results in a theory which is exclusionary toward trans individuals (the *trans inclusion problem*) and understanding gender as determined by an individual's gender identity results in a theory which is exclusionary toward disabled individuals (the *disability inclusion problem*), it begins to seem as though there is no suitably inclusive feminist metaphysical account of gender. This is the *gender inclusion problem*.

In this paper, I argue that it is possible to solve the gender inclusion problem by developing a theory of gender that incorporates aspects of both social-position accounts and identity accounts. According to the *structured* account I favor, an individual's gender is determined by their gender identity unless they lack such an identity, in which case it is determined by the way they are thought about and treated by other members of society. After introducing my structured account of gender and showing how it solves the gender inclusion problem, I compare my approach with some recent proposals due to Jenkins (2023), Barnes (2020), and others. I argue that the structured approach to gender offers a more direct—and therefore more convincing—solution to the gender inclusion problem than these alternatives. Finally, I explore inclusion problems that arise at the level of gendered language, showing that my structured account of gender can serve as the basis for a simple solution to them, as well.

1. Feminist metaphysics of gender

In this section and the next, I offer a brief overview of major themes in feminist metaphysics of gender and describe recent discussions of the gender inclusion problem.

What are the ambitions of feminist accounts of gender? As I will understand it, the feminist project in the metaphysics of gender aims to identify social categories suitable to do a certain kind of explanatory and emancipatory work. The fact that some individuals in contemporary Western society are expected to shave their legs is a social phenomenon, as is the fact that those individuals, considered as a group, have been (and continue to be) denied educational and professional opportunities.³ We want a descriptive account of the social categories—*gender* categories—which explain these kinds of social phenomena.⁴ When it comes to the emancipatory function of feminist accounts of gender, one widely shared hope is that a clear theoretical picture of these social categories will help to explain the historical evolution of the oppressive structures which generate and sustain them, and that this in turn will suggest present interventions to subvert those structures.

Haslanger (2000) provides an especially clear statement of this project. She writes that “the task is to develop [an account] of gender . . . that will be [an] effective [tool] in the fight against injustice,” and that for this reason an adequate account should help us “identify and explain persistent inequalities between females and males” and “track how gender . . . [is] implicated in a broad range of social phenomena extending beyond those that obviously concern sexual . . . difference” (2000, 36). Such an account will help us to explain the origin of gendered oppression in so far as “Prior to [historical] explanation it is valuable to provide clear conceptual categories to identify the phenomenon needing explanation, e.g., categories that identify the kind of injustice at issue and the groups subject to it.” (2000, 36).

Haslanger follows common practice among feminist philosophers in using the terms *woman* and *man* to refer to the social categories which play this explanatory and emancipatory role within feminist theory. On this subject, she writes:

The responsibility is ours to define [gender terms] for our purposes. . . . there is a stipulative element to the project: *this* is the phenomenon we need to be thinking about. Let the term in question refer to it. . . . Although the analyses I offer will point to existing social kinds (and this is no accident), I am not prepared to defend the claim that these social kinds are what our . . . gender talk is “really” about. My priority in this inquiry is not to capture what we do mean, but how we might usefully revise what we mean for certain theoretical and political purposes. (2000, 34)⁵

Here is one possible way of adding detail to my conception of the feminist project in the metaphysics of gender. We begin with what we are given: our pretheoretical or folk understanding of women and men. This enables us to observe the life experiences of various particular women and men and distill from our observations certain theoretically interesting generalizations (e.g. women are expected to shave their legs, women are denied educational and professional opportunities). Reflection on these generalizations enables us to grasp enough of the nature of patriarchy to realize that there is important emancipatory work to be done by changing the social conditions experienced by women.

To move from pretheory to theory, we identify a set of social phenomena which we want our theory of gender to help account for (call these the *explanatory desiderata*) and a set of practical goals which we want our theory of gender to help us achieve (call these the *emancipatory desiderata*). We then combine the explanatory and emancipatory desiderata into a reference-fixing description for our gender terms as follows: we associate our gender terms with social kinds according to the mapping that furnishes us with the metaphysical theory which achieves the best combination of explaining the explanatory desiderata and promoting the emancipatory desiderata. For example, a feminist metaphysician defending a social-position account of the categories *woman* and *man* might argue that, among possible accounts that can help us understand pretheoretical observations like our observation that women are denied educational and professional opportunities, mapping the term “woman” to the class of individuals thought about and treated by others in a certain (subordinating) way and the term “man” to the class of individuals thought about and treated by others in a certain other (privileging) way best promotes the practical project of resisting inequalitarian social structures.

The preceding suggestion is not the only way of understanding the feminist project in the metaphysics of gender, and I do not offer it as an interpretation of any particular feminist theorist. Nevertheless, I believe it is illuminating, and I will return to it at times in what follows as a way of clarifying the dialectic in certain debates about the nature of gender.

As I have described it, the feminist project in the metaphysics of gender is conceptually distinct from certain closely related feminist projects. For example, it is conceptually distinct from the revisionary linguistic project of advocating that we adjust the meanings of ordinary English words like “woman” and “man” so that they denote the feminist’s metaphysical gender categories.⁶ While many feminists are best read as engaging in both projects, there is no incoherence in seeking to develop the best feminist metaphysical theory of gender without also committing to the claim that ordinary English ought to be revised for feminist purposes. Though it is beyond the scope of my argument here to engage deeply with these related feminist projects, I consider the relationship between the metaphysics of gender and revisionary linguistic proposals in section 4 below.

In introducing social-position and identity accounts of gender and describing the gender inclusion problem above, I have used the terms *woman* and *man* to refer to gender categories. While this is more or less unavoidable, at least at the outset, for lack of

any widely accepted alternative terminology, it is worth emphasizing that these are technical terms which are not to be read with their standard English meanings. This point has an important methodological implication: intuitive judgments about the applicability of English words like “woman” and “man” in various actual and counterfactual circumstances do not constitute strong evidence either for or against feminist theories of gender.⁷ It is one thing to offer a theory of the ordinary meaning of the English word “woman”; it is quite another to offer a feminist metaphysical theory of gender categories.⁸ Indeed, one might offer a theory of the English word “woman” on which its meaning had little or nothing to do with the underlying metaphysical facts about gender. Such a theory is defended, for example, by Barnes (2020).⁹

Though it is coherent to use terms like “woman” and “man” in this stipulative way, the practice is not ideal in so far as it offers no simple way to signal whether one’s subject matter in claims about women is individuals falling into the extension of the ordinary English term or the feminist technical term. Different feminist authors make different choices here. For example, Barnes (2020, 714) reads Haslanger (2000) as using “woman” and “man” as technical terms. For her own part, however, Barnes uses them as ordinary language gender terms (2020, 725).¹⁰ And Jenkins (2023) uses the terms “woman” and “man” stipulatively to pick out social kinds which occupy a certain theoretical role. To avoid confusion, it is preferable to introduce a way of referring to the feminist metaphysician’s gender categories without using natural-language gender terms like “woman”. I will therefore adopt the terminology of referring to the category picked out by “woman” in its technical sense within feminist theory as “Gender 1” and to the category picked out by “man” in its technical sense within feminist theory as “Gender 2.”

2. The gender inclusion problem

We have seen that feminist accounts of gender can generally be classified as either social-position accounts or identity accounts. The social-position approach, which has roots in the work of Beauvoir (1953), holds that an individual’s gender is a property they possess in virtue of how they are perceived and treated by other members of society. According to social-position theorists, gender is determined neither by reproductive biology nor by an individual’s internal sense of oneself; instead, gender is that phenomenon whereby individuals are divided into social groups on the basis of what is *taken to be true* about their reproductive biology. The most influential recent social-position account of gender is due to Haslanger (2000). Haslanger holds that to belong to Gender 1 is to:

- (i) Be regularly and for the most part observed or imagined to have certain bodily features presumed to be evidence of a female’s biological role in reproduction,
- (ii) Be marked on the basis of these features within the dominant ideology of one’s society as someone who ought to occupy certain kinds of social positions that are in fact subordinate, and
- (iii) Be systematically subordinated along some dimension at least in part because of (i) and (ii).¹¹

In other words, for Haslanger, a woman in the sense relevant to feminist metaphysics is a person who has been assigned a certain kind of subordinate social role by others on the basis of their beliefs about her sex. This is a reductive view: it holds that facts about gender reduce to facts about how individuals are observed, imagined, and treated.¹²

Because social-position accounts of gender ground an individual's gender in the way they are thought about and treated by others, they do not admit the possibility that a person who is regularly thought about or treated as having male reproductive anatomy might belong to Gender 1. As Jenkins (2016) argues, this means that social-position accounts are subject to a trans inclusion problem. Jenkins asks us to consider the situations of trans women who do not publicly present as women, or who publicly present as women but whose presentation as women is not respected.¹³ Since these women fail to satisfy condition (i) of Haslanger's account, it predicts that they do not belong to Gender 1. But, Jenkins argues,

Failure to respect the gender identifications of trans people is a serious harm and is conceptually linked to forms of transphobic oppression and even violence. It follows from this that an important desideratum of a feminist analysis of gender concepts is that it respect these identifications by including trans people within the gender categories with which they identify and not including them within any categories with which they do not identify. (2016, 396)

Moreover,

if we overlook gender as identity, we will not be able to explain how oppression can operate through self-policing behavior even in the absence of external coercion. We will also find it difficult to understand how people can experience certain forms of gender-based oppression even if they are not classed as women. (2016, 415)

It is helpful here to interpret Jenkins through the lens of the feminist methodology described in section 1. As I read her, Jenkins is suggesting that respecting the gender identifications of trans people and ending transphobic oppression are legitimate feminist political aims, which should be included among the emancipatory desiderata on a feminist account of gender. Furthermore, she argues, there are certain recognizably gendered social phenomena, like self-policing, which should be included among the explanatory desiderata on a feminist account of gender. For this reason, she concludes, the account of gender that achieves the best combination of satisfying the feminist's explanatory and emancipatory desiderata will include trans people within the gender categories with which they identify. Since Haslanger's account does not have this feature, we should seek a different feminist account of gender.¹⁴

We can formulate the trans inclusion problem as follows:

(Trans Inclusion Problem): A feminist account of gender faces the trans inclusion problem just in case it fails to entail that all trans women belong to Gender 1 or that all trans men belong to Gender 2.

One influential response to the trans inclusion problem is to reject social-position accounts of gender in favor of accounts which tie gender exclusively to gender identity.¹⁵ Identity accounts of gender solve the trans inclusion problem by holding that an individual's gender is grounded in their own internal states as opposed to facts about how they are thought about and treated by others. According to these accounts, trans women who do not present as women, or who present as women but whose presentation is not respected, nevertheless identify with Gender 1 and thereby come to belong to Gender 1.

Metaphysical accounts of gender which tie it to gender identity thereby incur a responsibility to provide an explication of gender identity. In what follows, I will remain as neutral as possible concerning the correct theory of gender identity. For concreteness, however, it will be helpful to introduce Jenkins's influential norm-relevance account.¹⁶ For Jenkins, an individual's gender identity is grounded in their internal "map" for navigating the social world: "An individual has a female gender identity just in case that individual's internal 'map' is formed to guide someone classed as a woman through the social or material realities that are, in that context, characteristic of women as a class." (2016, 410). Jenkins later characterizes what it is to have such an internalized map in terms of what norms an individual takes to be relevant to their behavior: for example, individuals with a female gender identity understand the norm that women ought not have hairy legs to apply to them, whether or not they ultimately decide to behave in accordance with that norm (2016, 411).¹⁷ Given that Jenkins understands "classed as a woman" and "women as class" along the lines of social-position accounts of gender, this is a reductive view: it holds that facts about gender reduce to facts about whether the norms individuals take to be relevant to themselves are formed to guide individuals who are observed, imagined, and treated in a certain way.¹⁸

Despite their advantages, as Barnes (2020, 2022) points out, identity accounts are subject to their own inclusion problem: the disability inclusion problem. For it seems that some individuals—in particular, some intellectually disabled individuals—might experience recognizably gendered forms of oppression while lacking the sort of internal "map" required for having a gender identity. This means that identity accounts of gender will fail to classify them as belonging to any gender category. Barnes argues that this result is undesirable from a feminist perspective for two reasons. First, a theory of gender which fails to include in Gender 1 some intellectually disabled individuals who are socially positioned in the manner of women and who experience characteristically gendered forms of oppression is descriptively inadequate:

Cognitively disabled women are among the people most vulnerable to sexual abuse and rape. They are sterilized without their consent. They are routinely denied access to information about sexuality and birth control . . . They are even, on occasion, subjected to extensive medical procedures that will prevent the development of a typical adult female body . . .

Moreover, this kind of treatment isn't explained simply by how we treat people with cognitive disability. A substantial body of research suggests that there are striking gender differences in the treatment of cognitively disabled people. All cognitively disabled people are at increased risk of sexual abuse, but cognitively disabled women are especially vulnerable . . . (2022, 847)

And so I think we can only understand the experiences of cognitively disabled women if we recognize them as women. Likewise, we can only understand the full range of women's distinctive social experiences if we include cognitively disabled women. (2022, 849–50)

Second, such a theory fails to achieve some legitimate feminist political aims: "denying gender to cognitively disabled people further marginalizes them, even if we grant that

freedom from gender assignments might sometimes be a good thing.” (2022, 846). Barnes’s argument concerning disabled women thus parallels Jenkins’s argument concerning trans women: For Barnes, the account of gender that achieves the best combination of satisfying the feminist’s explanatory and emancipatory desiderata will include all disabled women in Gender 1. Since identity accounts do not have this feature, we must look elsewhere for an adequate feminist account of gender.¹⁹

We can formulate the disability inclusion problem as follows:

(Disability Inclusion Problem): A feminist account of gender faces the disability inclusion problem just in case it fails to entail that that all disabled women belong to Gender 1 or that all disabled men belong to Gender 2.

Neither the trans inclusion problem nor the disability inclusion problem is insoluble. As we have seen, identity accounts are designed to solve the trans inclusion problem, and social-position accounts are well equipped to solve the disability inclusion problem. Yet when the two problems are disjoined, it becomes less clear how to develop a feminist account of gender which adequately responds to them. If social-position accounts cannot escape the trans inclusion problem and identity accounts cannot escape the disability inclusion problem, perhaps no feminist account of gender can avoid being exclusionary. This more general problem, which I have called the gender inclusion problem, can be formulated more precisely as follows:

(Gender Inclusion Problem): A feminist account of gender faces the gender inclusion problem just in case it faces the trans inclusion problem or the disability inclusion problem.

Both social-position accounts and identity accounts—thus all major feminist accounts of gender—face the gender inclusion problem.²⁰

Of course, the gender inclusion problem is a problem only for those with a certain theoretical and methodological perspective: that of a feminist metaphysician who believes it is important to identify a single social kind which includes both all trans women and all disabled women. There are those who depart from this perspective in various ways. Some, for example, deny that it is important for feminist purposes to give any substantive theory of social gender kinds, while others hold that feminist purposes can be served by identifying a set of social gender kinds, some of which include all trans women and others of which include all disabled women.²¹

The fact that the gender inclusion problem only threatens those with certain theoretical and methodological starting points should not be taken to diminish the interest of a proposed solution for two reasons. First, the feminist metaphysical project in the context of which the problem arises has been and continues to be influential, so the question of whether it can be suitably inclusive is important in its own right. Second, as I argue in section 4 below, certain alternative theoretical and methodological perspectives that have risen to prominence in the context of worries about the gender inclusion problem are subject to their own worries—worries that strike me as less tractable than the gender inclusion problem, or which could only be satisfactorily addressed by adopting something structurally similar to the account I favor. Solving the gender inclusion problem is therefore indirectly relevant to our assessment of the attractiveness of these perspectives.

3. An inclusive theory

Discussions of the metaphysics of gender are often conducted under the assumption that social-position accounts and identity accounts are mutually exclusive in the sense that no single concept of gender can draw on both social position and identity. Jenkins (2016), for example, responds to the inclusion problem for Haslanger's social-position account by suggesting that we should acknowledge two distinct concepts of gender: one which works exclusively in a social-position way and one which works exclusively in an identity-based way.

Yet there is no reason why the two approaches might not be combined into a single account of gender. The trans inclusion problem for social-position accounts of gender suggests that when an individual has a developed gender identity, this identity determines their gender. At the same time, the disability inclusion problem for identity accounts of gender suggests that social position must be able to determine gender when an individual lacks a developed gender identity. Putting these observations together, we get the idea that, if an individual has a gender identity, this identity determines their gender, while if an individual lacks a gender identity, their gender is determined by how they are thought about and treated by others.

Following Barnes (2020), let us introduce the terms *feminized* and *masculinized* to describe the gendered ways in which different individuals are thought about and treated. Importing Haslanger's account of the relevant social relationships (but not construing it as a complete metaphysical story about gender), to be feminized is to:

- (i) Be regularly and for the most part observed or imagined to have certain bodily features presumed to be evidence of a female's biological role in reproduction,
- (ii) Be marked on the basis of these features within the dominant ideology of one's society as someone who ought to occupy certain kinds of social positions that are in fact subordinate, and
- (iii) Be systematically subordinated along some dimension at least in part because of (i) and (ii).

The same goes, *mutatis mutandis*, for being masculinized.

Then we can formulate an inclusive, structured metaphysical account of genders as follows:

(Inclusive Genders): If an individual has a gender identity, they belong to Gender 1 just in case they identify with Gender 1. If an individual lacks a gender identity, they belong to Gender 1 just in case they are feminized in a suitably broad range of contexts. The same is true, *mutatis mutandis*, for Gender 2.²²

What does it mean to say that an individual identifies with Gender 1? To say this is just to say that they instantiate the property other theorists have picked out using phrases like *identifies as a woman* or *has a female gender identity*. Though I remain as neutral as possible regarding the correct theory of gender identity, for present purposes I would be happy to adopt a version of Jenkins's internal map account of gender identity, described in section 2, and say that what it is for an individual to identify with Gender 1 is for that individual to have an internal map that is formed to guide feminized individuals through the social or material realities that are characteristic of feminized individuals as a class. The same holds, *mutatis mutandis*, for Gender 2. Thus (Inclusive Genders) is a reductive

view: it holds that facts about gender reduce to (i) facts about how individuals are observed, imagined, and treated, and (ii) facts about whether the norms individuals take to be relevant to themselves are formed to guide individuals who are observed, imagined, and treated in a certain way.²³

(Inclusive Genders) solves both the trans inclusion problem and the disability inclusion problem. Since part of what it is to be a trans woman is to identify with Gender 1, it predicts that trans women belong to Gender 1 in virtue of their gender identity regardless of whether they are feminized in any context—and the same goes, *mutatis mutandis*, for trans men. With respect to the disability inclusion problem, (Inclusive Genders) predicts that, even if they lack a gender identity, intellectually disabled women can belong to Gender 1 in virtue of being feminized in a suitably broad range of contexts.

There is thus much to be said for (Inclusive Genders). Some metaphysicians might hesitate to endorse it, however, because it may strike them as disjunctive, disunified, or gerrymandered. After all, (Inclusive Genders) holds that there are two quite different ways in which an individual's gender might be grounded in more fundamental facts about the world: it might be grounded in facts about their gender identity, or it might be grounded in facts about their social position.

There are several points to be made in defense of (Inclusive Genders) here. First, the feminist project in the metaphysics of gender, as I have articulated it in section 1, is not a purely descriptive project. The search for social gender categories is constrained by both our explanatory and our emancipatory desiderata. Given the constraints imposed by our emancipatory desiderata, it is not obvious that we should *ex ante* expect the social gender categories that achieve the best combination of explaining and emancipating to be particularly simple. When the world and our emancipatory aims are complicated, so too must be our metaphysical theories.

This is not to deny that simplicity is a theoretical virtue. I think the friend of (Inclusive Genders) should concede that, *ceteris paribus*, less disjunctive metaphysical proposals are to be preferred to more disjunctive ones. But the relevant simplicity comparison is between rival proposals that achieve the best combination of satisfying our explanatory and emancipatory desiderata, and it may turn out that none of these proposals is particularly simple. This is important, since it follows that the greater simplicity of alternative metaphysical theories of gender categories constitutes an objection to (Inclusive Genders) only if those theories can be shown to achieve an equally good combination of satisfying our explanatory and emancipatory desiderata. As I argue in the next section, existing unified metaphysical accounts of gender are subject to the gender inclusion problem, and this means that *ceteris* are not *paribus* when we compare them with (Inclusive Genders): because they fare poorly with respect to the feminist's emancipatory desiderata, they are not "in the running" to compete with (Inclusive Genders) on simplicity.

Second, it is worth noting that disjunctive, disunified, or gerrymandered social properties are not particularly uncommon. Consider the fact that a couple is married. In jurisdictions that recognize common-law marriage, this fact might be grounded in either of two quite different ways: on the one hand, the couple might have taken part in a legally recognized marriage ceremony and obtained a marriage license; on the other, they might simply have lived together in the manner of a married couple for long enough. A metaphysical account of the social property of being married as it is constructed in such jurisdictions would need to be disjunctive in order to issue correct theoretical verdicts. For the same reason, it would be inappropriate to rule out (Inclusive Genders) on general methodological grounds.²⁴

Third, I argue in section 5 below that the additional complexity built into (Inclusive Genders) allows us to endorse an attractively simple account of how we could implement a trans and disability inclusive linguistic practice. So when one considers the combined theoretical package of a feminist metaphysics and a feminist philosophy of language, the complexity of (Inclusive Genders) may be offset by gains in simplicity elsewhere.²⁵

4. An alternative approach?

Let's take stock. So far, I have described the gender inclusion problem and shown why both social-position accounts of gender and identity accounts of gender seem unable to solve it. I have also presented a novel structured account of gender categories—(Inclusive Genders)—and shown how it is able to solve the gender inclusion problem.

Of course, my proposed solution to the gender inclusion problem does not exist in a vacuum. There are a number of other approaches to inclusion worries which have recently been explored in the literature, including a pluralist proposal advocated by Jenkins (2016, 2023), a disjunctive account developed by Bettcher (2013), and a pair of structurally similar recent proposals by Barnes (2020) and Antony (2020). In this section, I compare these strategies to the solution offered by (Inclusive Genders) and give my reasons for preferring the latter.

The most prominent recent advocate of pluralism about gender categories has been Jenkins (2016, 2023), who argues that both the concept of gender as social position and the concept of gender as identity play important explanatory roles in feminist theory, and that we should regard neither as primary. As we have seen above, Jenkins is sensitive to issues of trans inclusion, and her conceptual pluralism is in part motivated by the thought that no trans-exclusionary account of gender could satisfy the feminist metaphysician's explanatory and emancipatory desiderata.

As I read Jenkins, her pluralism can be thought of as a kind of ambiguity thesis: there is more than one set of social kinds that answers to the feminist metaphysician's reference-fixing description by best satisfying her explanatory and emancipatory desiderata. So it does not make sense to speak of gender, or of gender categories like Gender 1 and Gender 2, *simpliciter*. Instead, we must always specify whether we mean identity-gender (and, correspondingly identity-Gender 1 and identity-Gender 2) or social-position-gender (and, correspondingly, social-position-Gender 1 and social-position-Gender 2). For some feminist purposes, it will be most helpful to think in terms of identity-gender; for others, it will be most helpful to think in terms of social-position-gender.

In her earlier work, Jenkins (2016) proposes to ensure that her pluralist account of gender is trans-inclusive by arguing that, while identity-gender and social-position-gender are equally important for feminist purposes, the word "woman" ought to be used only in an identity-based sense (see section 5 for further discussion). More recently, Jenkins develops an account of what she calls *ontic injustice*—"the very fact that an individual is a member of a social kind can be wrongful, because the constraints and enablements that partly constitute kind membership can fall short of what is morally owed to the individual who is subjected to them" (2023, 5)—and argues that recognizing the possibility of ontic injustice motivates a rejection of what she calls *the Ontology-First approach*: "the assumption . . . that settling questions about the current ontology of gender kinds will automatically determine what shape our gendered social practices ought to take" (2023, 201). Rejecting the Ontology-First approach is conducive to trans

liberation, she argues, because it allows us to bypass debates about who is “really” a woman in constructing or reconstructing our gendered social practices.

Where Jenkins argues that we need multiple gender concepts to do the explanatory and emancipatory work of feminist metaphysics, I have argued that the single concept articulated in (Inclusive Genders) is sufficient. And where Jenkins argues that the best way to promote trans inclusion is to reject the Ontology-First approach, I have suggested that the best way to promote trans inclusion is to defend a trans-inclusive account of gender categories—and thus implicitly endorsed the Ontology-First approach. Why prefer my approach to inclusion problems to Jenkins’s pluralism? By way of answering this question, I will respond to a worry that my structured account is less explanatory than Jenkins’s pluralist approach, describe why I am not convinced by Jenkins’s main argument against the Ontology-First approach, and then present a problem for the pluralist which I believe motivates seeking an alternative theory.

A first concern is whether my structured account of gender has the resources to do the same explanatory work as a pluralist approach to gender. In particular, one might worry that the natures of the trans exclusion problem and the disability exclusion problem are quite different: the disability exclusion problem is about explaining how disabled women can be oppressed qua women, while the trans exclusion problem is about recognizing individuals’ right of self-determination with respect to gender category membership. So perhaps a theory which posits a single concept of gender will unavoidably flatten the phenomena to be explained by failing to recognize important distinctions. For example, how might a structured account of gender like (Inclusive Genders) explain how some trans men can be subject to misogyny?

I think my structured account of gender fares just as well as a pluralist approach when it comes to satisfying the feminist’s explanatory desiderata. After all, endorsing (Inclusive Genders) does not require us to do away with the concept of feminization, or of gender identity—in fact, it requires us to keep them, since they feature among the determinants of gender category membership. Given this fact, we can explain the difference between the experiences of oppression of certain disabled women and certain trans women by saying that the disabled women are oppressed primarily in that they are feminized and feminized individuals are marked for subordination by social power structures, while the trans women are oppressed primarily in that the fact that they belong to Gender 1 is not recognized by those around them. Trans men may be oppressed in both these ways simultaneously: because the fact that they belong to Gender 2 is not recognized by those around them, they may be feminized and therefore marked for subordination. So (Inclusive Genders) has the resources to make sense of a wide range of experiences of oppression without flattening the phenomena.²⁶

A second concern is whether theorizing about gender in a way that embraces the Ontology-First approach is necessarily counterproductive from the perspective of trans inclusion. Jenkins offers an argument that this is the case:

As we have seen, the Ontology-First approach directs us to look for kinds that we can point to and say, “Look, that’s what gender *really* is—so that’s what our social practices should track”. For those who want to justify trans-inclusive social practices, a natural choice for which kinds to highlight in this way would be gender identity kinds . . .

. . . the Ontology-First approach funnels proponents of trans-inclusive social practices towards making the claim that gender identity kinds can do the

explanatory work that we've historically asked the idea of "gender" to do for us. This claim, as I have demonstrated in Chapters 4 and 5, is false . . . [in addition to gender identity,] oppression also has a lot to do with how other people treat us and how we are positioned within social structures, and this might proceed in ways that don't line up with a person's gender identity . . . If I am right, then the Ontology-First approach funnels proponents of trans-inclusive social practices towards trying to defend a claim that is false.

I agree with Jenkins that the Ontology-First approach would be problematic if it required feminists committed to trans inclusion to defend a false claim. But I think the idea that in order to respect the identifications of trans individuals we must embrace an account of gender categories that is insensitive to social position is mistaken: we can instead opt for a structured account like (Inclusive Genders), which, as I have argued above, has the resources to explain the oppression individuals face in virtue of how they are positioned in social structures.

So I do not think there are strong negative arguments against my structured approach as compared to a pluralist approach. I do, however, think there are important negative arguments against a pluralist approach. For there is an important sense in which pluralist approaches fail to solve the gender inclusion problem: by abandoning gender *simpliciter* in favor of identity-gender and social-position-gender, they embrace a set of gender concepts all of which are exclusionary. Do all trans women belong to Gender 1? Well, all trans women belong to identity-Gender 1, but not all trans women belong to social-position-Gender 1. And the same problem arises, *mutatis mutandis*, for the question of whether all disabled women belong to Gender 1. To bring this point home, consider that friends of pluralist approaches are committed to endorsing claims like:

- (i) There is a perfectly legitimate feminist sense of gender in which many trans women do not share a gender with cis women.
- (ii) There is a perfectly legitimate feminist sense of gender in which many disabled women do not share a gender with nondisabled women.
- (iii) There is no legitimate feminist sense of gender in which all trans women and all disabled women share a gender.

And if we imagine some particular non-feminized trans woman Ada and some particular severely cognitively disabled woman Beth, friends of pluralist approaches are committed to endorsing:

- (iv) There is no legitimate feminist sense of gender in which Ada and Beth share a gender.

These consequences of pluralist theories of gender strike me as undesirable if we take on board that trans- and disability-inclusivity are emancipatory desiderata on a feminist theory of gender. As Jenkins herself remarks, "It seems as though what some people really want here is an account of the ontology of gender kinds that . . . guarantees that the only correct claim that can be made about [a trans woman's] membership in a gender kind is that she is a woman." (2023, 195). Unlike pluralist approaches, my structured account of gender does have the feature Jenkins describes here, which I think it is reasonable for a trans-inclusive feminist to want. For this reason, my inclination is to

say that, while appealing to multiple gender concepts to explain various gender-related phenomena may be a viable feminist fallback position, it is preferable to explain those phenomena using a single structured gender concept.

I turn next to Bettcher (2013), who develops an account of the meaning of “woman” as it is used in trans-inclusive communities that suggests a metaphysics of gender categories which may at first seem quite similar to my structured approach. Bettcher suggests that the meaning of the word “woman” is modified in trans-inclusive contexts using a “two-step process”: first, the extension of the expression “trans woman” is specified without reference to the extension of “woman” (Bettcher suggests that a person falls in the extension of “trans woman” just in case “she was assigned to the male sex at birth, currently lives as a woman, and self-identifies as a trans woman (or as a woman)”; second, the extension of “woman” is fixed so that it includes “*both* trans and non-trans women (where [a] ‘non-trans woman’ is a person who counts as a woman but does not count as a trans woman)” (2013, 241; emphasis in original).

As written, Bettcher’s account appears to involve a circularity in specifying the extension of “woman,” since in order to know which individuals fall in the extension of “non-trans woman” we must know which individuals fall in the extension of “woman.” But I think Bettcher’s intention here is relatively clear: the extension of “non-trans woman” is fixed using some existing account of the extension of “woman”—in the paper, Bettcher considers sex-based accounts and family resemblance accounts.²⁷ Letting “Property W” stand in for whatever property is held to determine membership in the extension of “non-trans woman” and “Property M” stand in for whatever property is held to determine membership in the extension of “non-trans man,” we can formulate Bettcher’s proposal as a metaphysical claim about gender category membership as follows:

(Disjunctive Genders): An individual belongs to Gender 1 just in case either they are a trans woman or they possess Property W. An individual belongs to Gender 2 just in case either they are a trans man or they possess Property M.

(Disjunctive Genders) is similar to (Inclusive Genders) in that it holds that gender category membership is in some cases determined by gender identity and in other cases determined in another way. For this reason, (Disjunctive Genders) goes some way toward solving the trans inclusion problem. There are, however, two worries for (Disjunctive Genders) which do not arise for (Inclusive Genders).

First, the options for Property W that Bettcher considers give rise to their own kinds of inclusion problems. On a sex-based account of Property W, for example, non-trans intersex people who are socially positioned as women will fail to belong to Gender 1 according to (Disjunctive Genders): their gender identities will not qualify them for membership in Gender 1 because they are not trans. Similarly, on a family resemblance account, some gender nonconforming non-trans female individuals who are socially positioned as women will fail to belong to Gender 1. Because (Inclusive Genders) grounds gender category membership in social position rather than these alternatives, it avoids these problems.

Second, (Disjunctive Genders) makes inappropriate predictions about the gender category membership of trans individuals. Consider the kind of case that initially motivated Jenkins to advocate for an identity-based gender concept: the case of a trans woman whose gender identity is not respected by those around her. (Disjunctive Genders) counts such an individual as a member of Gender 1 in virtue of the fact that

she is a trans woman. But, since it is a simple disjunctive account, it also counts her as a member of Gender 2 if she possesses Property M. This means that whether we choose sex, family resemblance, or social position as our Property M, there will be cases where (Disjunctive Genders) predicts that a trans woman is also a man. To solve this problem, we would need to move away from the disjunctive nature of (Disjunctive Genders) and toward a structured view according to which gender identity determines gender category membership whenever it exists—that is, toward (Inclusive Genders).

Finally, I consider the approaches of Barnes (2020) and Antony (2020), who seek to avoid inclusion problems by offering no (or at least no substantive) metaphysical account of gender categories.

In developing a theory of the relationship between gender and gendered language, Barnes seeks to shield Haslanger's social-position account of gender from worries about trans inclusion by declining to construe it as a theory either of gender categories or of natural-language gender terms. For Barnes, the meanings of gender terms in natural language are a poor guide to the underlying metaphysics of gender. Once this is realized, she argues, it becomes possible to defend a Haslangerian social-position metaphysics from the trans inclusion problem by taking care not to interpret it as a semantic account of the word "woman" or an attempt to give "the full story about gender" (2020, 715). Instead, we can understand Haslanger's account as grounding the existence of gender as a social phenomenon in facts about the ways individuals are treated on the basis of what is observed or imagined about their bodies. While the social reality of feminizing and masculinizing cannot itself be the whole story about gender because it tells us nothing about the complex phenomena of gender identity and gender expression, Barnes maintains that "being masculinized or feminized in a context [is] the basic social reality that explains our complex social experience of gender" (2020, 715).

There is much to admire about Barnes's strategy for insulating social-position accounts of gender from the trans inclusion problem. But it is best understood as way of sidestepping rather than solving the gender inclusion problem. This becomes clear when we consider the distinction between offering a metaphysical account of gender, on the one hand, and offering a metaphysical account of gender categories, on the other.

The existence of gender is undoubtedly a genuine social phenomenon, which must be explained by our total metaphysical theory. Individuals with reproductively different bodies are assigned different social roles on this basis, and these social roles shape and constrain their lives in myriad ways. Gender exists because we individually and collectively recognize and enforce this connection between bodies and social roles.

The existence of gender categories is also a real social phenomenon. Whereas gender is the phenomenon grounded in the assignment of social significance to reproductive difference, gender categories are the particular social categories into which individuals are sorted. Gender is a determinable, while gender categories are determinates thereof.

A metaphysical account of the existence of gender is not itself a metaphysical account of gender categories. The fact that individuals with reproductively different bodies are assigned different social roles does not itself tell us how many such social roles there are or what their membership conditions might be. It is plausible that gender as a phenomenon could not exist without there being at least two genders—but we could have constructed social reality such that there were more than two gender categories.²⁸ And while the existence of genders is grounded in the assignment of social roles on the basis of reproductive difference, it does not follow that the membership conditions of particular gender categories turn in any simple way on biology or social position.

As Barnes puts it, “The underlying structure of gender is binary, but that needn’t mean that there are only two genders” (2020, 717).

It follows that, if we are interested in constructing a trans-inclusive feminist theory of gender categories, Barnes’s strategy is incomplete by design. It does not tell us whether trans women and cis women belong to the same gender category, or whether cognitively disabled women and nondisabled women belong to the same gender category. In so far as the gender inclusion problem requires a feminist metaphysical account of gender to entail that trans women belong to Gender 1, then, Barnes’s strategy does not solve the gender inclusion problem.²⁹

Like Barnes, Antony (2020) seeks to explain the existence of gender without providing a substantive account of gender categories and their membership conditions. For Antony, however, “human sex difference is the *material ground* of systems of gender” in that “gender systems serve to enable male control of female reproductive capacities” (2020, 529; emphasis in original). This claim about the material ground of systems of gender does not, however, entail that membership in gender categories is determined by sex. Instead, Antony rejects the idea that feminist purposes require any substantive theory of gender categories or their membership conditions:

I contend that, whatever the inherent philosophical interest of such questions as “what is it to be a woman,” there is no practical or political need to answer them . . . we can (and should) regard ‘woman’ as a semantically primitive term, and WOMAN as a primitive concept, both having as their extension members of the primitive social kind, woman. (2020, 532)

Antony suggests that her deflationary account is preferable to social-position accounts because it does not entail that non-passing trans individuals fail to belong to the genders with which they identify (2020, 542). While this is a first step toward solving the trans exclusion problem, it is not itself a solution, since Antony’s account also does not entail that non-passing trans individuals *do* belong to the genders with which they identify. Indeed, Antony’s deflationary account does not entail anything substantive about who belongs to which gender category. So, while it differs from gender concept pluralism in that it does not entail that sentences (i–iv) above are true, it is similar to Barnes’s approach in that it fails to entail that they are false. In so far as we take being trans and disability inclusive as desiderata on a feminist account of gender, then, we have reason to prefer (Inclusive Genders) to Antony’s approach.

5. Linguistic inclusion

So far, I have been concerned with inclusion problems arising at the level of the metaphysics of gender. But many feminists have also worried about inclusion problems arising at the level of gendered language. After all, the same emancipatory considerations that motivate trans and disability inclusive gender categories plausibly also motivate adopting trans and disability inclusive linguistic practices. For some feminists who have defended non-inclusive gender concepts, linguistic inclusion has even come to be regarded as more central to the feminist project than metaphysical inclusion. In this section, I formulate the inclusion problems that arise at the level of gendered language, assess some recent approaches to solving them, and then show how (Inclusive Genders) points the way to an especially simple and elegant solution.

Whereas inclusion problems at the level of metaphysics can be stated straightforwardly in terms of conditions on which individuals are grouped together into gender categories, formulating inclusion problems at the level of gendered language is more complicated. As in the case of metaphysical inclusion, the project of linguistic inclusion is driven in part by emancipatory considerations. But it is not immediately clear what conditions a linguistic practice must satisfy in order to be suitably emancipatory. Must it be such that certain sentences express truths or falsehoods? Or that certain utterances are assertible or unassertible? And is the project to describe a possible inclusive linguistic practice and recommend its adoption, or is it to establish that our actual linguistic practice is inclusive?

To begin, consider the question of what an account of the meanings of gender terms says about their extensions. For example, does a theory of the meaning of “woman” include both all trans women and all disabled women in the extension of that expression?³⁰ One might hope to formulate inclusion problems at the level of gendered language in terms of the extensions of gender terms. An immediate difficulty with this approach, however, is that questions of extension cannot be applied directly to theories which posit linguistic context sensitivity in the meanings of gender terms. According to such theories, gender terms may express different properties (and therefore have different extensions) when they are uttered in different contexts. A context-sensitive account of “woman,” for example, might hold that in some contexts it expresses a property possessed by all trans women but not all disabled women, while in other contexts it expresses a property possessed by all disabled women but not all trans women.³¹

We might respond to this difficulty by holding that a theory of gender terms fails to be suitably inclusive if it fails to include all trans and disabled women in the extension of “woman” and all trans and disabled men in the extension of “man” in every context. But the resulting criterion for assessing semantic theories would arguably be simplistic. Theories of meaning interact in complex ways with pragmatic principles governing felicitous assertion. Sentences which are true according to a semantic theory need not always be felicitously assertible (perhaps, for example, they are misleading), and sentences which are felicitously assertible need not always be true (perhaps, for example, they are polite, or close enough to true for conversational purposes, or figurative). Friends of a context-sensitive theory might hold that sentences like “Some trans women are not women” express truths in certain contexts but, even in those contexts, remain unassertible.³²

Given these complications, how might we understand inclusion problems arising at the level of gendered language? Here I will make two suggestions. First, I suggest that the key concern of feminists advocating for linguistic inclusion is to explain why it is generally not acceptable to use trans or disability exclusionary language. This goal might be achieved either by showing that our actual linguistic practice is such that trans and disability exclusionary language is generally semantically or pragmatically defective or by showing that we ought to adopt a linguistic practice which is suitably inclusive. Second, I suggest that the notion of trans and disability inclusion which feminists have had in mind be explicated in terms of whether a linguistic practice secures a positive enough status for the sentences “All trans women are women” and “All disabled women are women” (and the corresponding sentences involving men), and whether it adequately explains what is objectionable about the sentences “Trans women are not women” and “Disabled women are not women” (and the corresponding sentences involving men).

These suggestions are supported by a variety of remarks by feminist philosophers discussing gendered language. In her early work, for example, Jenkins argues that the

English word “woman” should be used “exclusively for the concept of *having a female gender identity*” because “an unintended effect of using ‘woman’ to refer to gender as class (even if it were also used to refer to gender as identity) would be the reinforcement and perpetuation of the existing marginalization of trans women within feminism” (2016, 417). I read Jenkins as arguing that, given the emancipatory goals of trans-inclusive feminism, we ought to adopt a linguistic practice associating the word “woman” with an identity-based gender concept, because if we do not we will not be able to secure a positive enough status for the sentence “All trans women are women” or explain what is objectionable about the sentence “Trans women are not women” (because there will be readings of these sentences on which they express truths).

Similarly, while Barnes disagrees with Jenkins in holding that we should “[allow] for flexibility and mutability in the way we use gender terms,” she argues that “we should, whenever possible, treat the sincere self-ascription of gender terms as true,” and that sentences like “Trans women aren’t women,” even if they express truths in some contexts, tend to pragmatically communicate false propositions about “the basic social reality of gender” and are therefore “incorrect and inappropriate thing[s] to say” (2020, 722). I read Barnes as arguing that our actual linguistic practice is such that pragmatic considerations assign a positive status to “All trans women are women” and a negative status to “Trans women are not women.”

In more recent work, Jenkins largely follows Barnes’s approach to linguistic inclusion problems, arguing that while “it would be a problem for my account, or at least a significant drawback, if it did sanction [a misgendering] utterance,” it does not do so because such utterances may be pragmatically defective (2023, 193). Even if they are neither semantically nor pragmatically defective, Jenkins adds, misgendering utterances will run afoul of “a final dimension of analysis . . . which we can think of as *prescriptive*: How should we be aiming to use words? What collective linguistic practices should we be fostering?” (2023, 194; emphasis in original). Here I read Jenkins as combining her earlier approach with Barnes’s: even if Barnes’s pragmatic claim about our actual linguistic practice turns out to be false in some cases, the feminist interest in gendered language is able to fall back on the idea that we *ought* to use gender terms in a trans-inclusive way.

We can formulate the trans inclusion problem for gendered language, as it has been understood by trans-inclusive feminists, as follows:

(Gendered Language Trans Inclusion Problem): An account of our linguistic practice involving gender terms faces the gendered language trans inclusion problem just in case (i) it fails to secure a positive enough status for the sentences “All trans women are women” and “All trans men are men,” or (ii) it fails to secure an adequate explanation of what is objectionable about asserting “Trans women are not women” and “Trans men are not men.”³³

Applying the criteria articulated in (Gendered Language Trans Inclusion Problem) involves making judgments about whether a given theory secures a positive enough status for trans-inclusive sentences and whether it secures an adequate enough explanation of what is objectionable about trans-exclusionary language. In some cases, these judgments will be easy: a theory which holds that pre-transition or non-passing trans women are not in the extension of “woman” in any context, and that there are not and ought not be any pragmatic or other reasons not to say so, clearly faces the gendered language trans inclusion problem. In other cases, however, the judgments may be less

clear, and there may be room for debate concerning whether a given theory faces the gendered language trans inclusion problem.³⁴

Both Jenkins and Barnes seek to solve the gendered language inclusion problem without suggesting a linguistic practice which classifies all trans and disabled women in the extension of “woman” in all contexts. On Barnes’s account of gendered language, for example, there will be many contexts in which “woman” expresses a trans-exclusionary or disability-exclusionary meaning. In such contexts, sentences like “Trans women are not women” and “Severely cognitively disabled women are not women” will express truths. As we have seen, Barnes’s response to this fact about her theory of gendered language is to turn to the false propositions pragmatically communicated by utterances of such sentences to explain what is objectionable about them. One might reasonably worry that this pragmatic strategy cannot succeed in sufficient generality: it is in the nature of pragmatic phenomena like implicatures, after all, that they can be cancelled. Even if the strategy does succeed, some may feel that it would be better for a theory not to have to turn to pragmatics to nullify the trans and disability exclusionary elements of its semantics. Similarly, Jenkins (2016) seeks to solve the gendered language trans inclusion problem by arguing that the word “woman” should express a gender-identity based property, and Jenkins (2023) offers this normative suggestion as a fallback position. But this normative approach to the gendered language trans inclusion problem comes at the cost of giving up on disability-inclusive gendered language, since some cognitively disabled women will not count as women according to the linguistic practice Jenkins recommends.

Happily, (Inclusive Genders) enables us to offer meanings for “woman” and “man” that solve the gendered language inclusion problem. We can simply hold that, in every context, “woman” ought to express the property *adult belonging to Gender 1* and “man” ought to express the property *adult belonging to Gender 2*. That is, we can offer the following semantic proposal as a description of a linguistic practice which should be adopted in light of the feminist’s emancipatory goals:

(Inclusive Extensions): An individual falls into the extension of “woman” if and only if they are an adult belonging to Gender 1, and an individual falls into the extension of “man” if and only if they are an adult belonging to Gender 2.³⁵

Because (Inclusive Genders) entails that all trans women belong to Gender 1 and all trans men belong to Gender 2, (Inclusive Extensions) is not subject to the gendered language trans inclusion problem. And because (Inclusive Genders) entails that the intellectually disabled women and men at issue in discussions of the disability inclusion problem belong to Gender 1 and Gender 2, respectively, (Inclusive Extensions) is not subject to the gendered language disability inclusion problem.³⁶

6. Conclusion

Theorists who hold that an individual’s gender is wholly determined by their social position struggle to accommodate the fact that all trans women are women and all trans men are men. Theorists who hold that an individual’s gender is wholly determined by their gender identity, on the other hand, struggle to accommodate the fact that all disabled women are women and all disabled men are men. In response to these problems, I have argued that it is possible to be inclusive by adopting a structured

metaphysics of gender sensitive to both individuals' gender identities and their social positions.

I have proposed a particular structured metaphysical theory of gender categories—(Inclusive Genders)—according to which an individual's gender identity determines their gender whenever this is possible, while facts about social position “fill in the gaps” where no gender identity is present. And I have shown how my structured theory of gender suggests a simple semantics for natural-language gender terms that avoids exclusion worries at the level of gendered language. My proposal fares well in comparison to other recent work seeking to respond to inclusion problems in the metaphysics of gender.

Acknowledgements. I am grateful to Carolina Flores and Isabel Uriagereka Herburger, as well as audiences at the Rutgers Feminist Philosophy Reading Group and the 2024 Swedish Congress of Philosophy (Filosofidagarna), for helpful comments and discussion.

Notes

1 This generalization is not without exceptions. Butler (1990, 1993), e.g., suggests that sex is, in some sense, a social phenomenon like gender (see Åsta 2011, 2018, for discussion of how to interpret this claim). And “gender-critical” feminists maintain that gender is determined in a direct way by sex, so that the two phenomena are not modally distinct (see, e.g., Lawford-Smith 2022, 92).

2 Identical exclusion worries arise when considering trans and disabled men.

3 Nothing about the methodology I describe here depends on these examples in particular. The example of leg-shaving expectations is an established one, discussed by Jenkins (2016) among others; the denial of educational and professional opportunities to women is one of the historically most significant manifestations of patriarchy. The list of gendered social phenomena which the feminist metaphysician is interested in explaining will naturally extend beyond these two examples. It should also be noted that the social phenomena which the feminist metaphysician aims to explain by appealing to her theory of gender categories need not be *wholly* explained by those categories. Because women did not always experience gendered pressure to shave their legs, the explanation for the current state of affairs must involve historical factors other than gender. Leg-shaving expectations are nevertheless a recognizably gendered social phenomenon.

4 Following Bernstein (2020, 325), it is helpful to distinguish between *gender*, which is a determinable, and particular *gender categories*, which are determinates thereof. More on this distinction in section 4.

5 Along similar lines, Mikkola (2022) describes how feminists have “appropriated the term ‘gender’” and used it such that “‘gender’ denotes women and men depending on *social* factors (social role, position, behaviour or identity)” with the hope that “since gender is social, it is . . . mutable and alterable by political and social reform that would ultimately bring an end to women’s subordination.”

6 As Barnes remarks, this conceptual distinction has not always been clearly emphasized: “there’s some confusion in the relevant literature over whether the target of analysis should be ordinary language terms like ‘woman’, our concept of *woman*, or perhaps a social category *being a woman*. Indeed, sometimes the target seems to be all three. And typically where focus is on the latter two the idea is that this is the underlying social reality which our use of terms like ‘woman’ ought to track.” (2020, 724).

7 This is not to say that they constitute no evidence at all. A feminist theory of gender categories might reasonably be criticized, for example, if it predicted that there was no statistical correlation between an individual’s gender category membership and whether they fell into the extension of the ordinary English terms “woman” or “man.” The point is that the feminist metaphysician is not threatened by local mismatches between facts about the extensions of English gender terms and facts about membership in the social categories she identifies as genders.

8 For example, while Byrne (2020) has recently offered a number of arguments for the conclusion that women are adult human females, I take it that he is engaged in the project of investigating the property expressed by the ordinary English gender term “woman” rather than the project of developing a metaphysical theory of gender which best satisfies the feminist’s explanatory and emancipatory desiderata. For further discussion of Byrne’s arguments, see Dembroff (2021), Heartsilver (2021), and Byrne (2021, 2022).

9 This methodological point is sometimes presented as though it follows more or less directly from the distinction between metaphysics and semantics. For example, Barnes writes that “The project of developing a philosophical theory of gender can and should come apart from the project of giving definitions or truth conditions for sentences involving our gender terms” (2020, 704). While this can be a helpful way to conceptualize the point, it may also be confusing in the following sense: unless one is prepared to reject the idea that English words like “woman” and “man” express properties and disquotational formulas like “ x is a woman” is true just in case x is a woman” express truths, one must concede that investigating the meanings of natural language gender terms is a way of investigating certain metaphysical features of the world.

To clarify, I understand the relationship between semantics and metaphysics in the context of feminist accounts of gender as follows. In addition to the ordinary English gender terms ‘woman’ and ‘man’, the feminist metaphysician introduces homonymous technical terms ‘woman’ and ‘man’ using something like the method described above. She endorses disquotational principles in full generality, so that both the ordinary English ‘woman’ and ‘woman’ in the technical sense express properties. However, she incurs no commitment to thinking that the properties expressed in various contexts by ordinary English gender terms play the explanatory or emancipatory roles in which she is interested *qua* feminist metaphysician.

10 This is why she makes claims like “The task of giving an explanation of the social reality of gender can and should come apart from the task of saying what it is to be a woman” (2020, 724). See also (2020, 725): “I’m going to primarily focus on [ordinary language] gender terms like ‘woman.’”

11 Haslanger’s account of what it is to belong to Gender 2 is structurally similar, though the relevant bodily features must be presumed to be evidence of male reproductive biology rather than female reproductive biology, and the social position which an individual is marked to occupy on the basis of these features must be dominant rather than subordinate.

12 Note that, while the idea that gender categories are constitutively tied to hierarchies of privilege and subordination plays an important role in Haslanger’s overall theoretical framework, it is conceptually independent from the dialectic surrounding problems of inclusion. Other social-position theorists offer accounts which are not constitutively tied to hierarchy in this way. According to Ásta (2011, 2013, 2018), for example, an individual’s gender is conferred on them in a context by others in a way that attempts to track some base property (e.g., their role in biological reproduction). Along similar lines, Witt (2011) holds that gender is “the principle of normative unity” (79) for human persons, where “whether an individual is a woman or a man is fixed by the reproductive role that individual is recognized by others to perform” (29). Neither Ásta nor Witt takes gender categories to be essentially hierarchical.

13 How should we understand the extension of “trans women” as it occurs in formulations of the trans inclusion problem? We cannot simply say that trans women are women who are trans, or that they are trans individuals who belong to Gender 1, since to do so would be to make the question of whether trans women belong to Gender 1 definitional rather than substantive. Instead, I follow Jenkins (2016, 396) in holding that trans women are “[adult] people categorized as male at birth who later come to identify as women”—or, in my preferred terminology, adult people who were male-bodied at birth and identify with Gender 1. Given the strategy described below for reducing facts about identifying with Gender 1 to facts about whether the norms an individual takes to be relevant to themselves are formed to guide individuals who are observed, imagined, and treated in a certain way, this characterization of the extension of “trans women” allows us to understand arguments for trans-inclusive theories of gender as substantive arguments for developing a metaphysical theory that includes a particular group of individuals, identifiable independently of their gender category membership, in Gender 1. The same strategy can be applied, *mutatis mutandis*, for “trans men,” “trans girls,” “trans boys,” and “trans individuals.”

14 Note that the same difficulty arises for the accounts of Ásta and Witt: a trans woman who does not present as a woman might not have the appropriate gender conferred on her by others, and she might not be recognized by others to perform a female reproductive role. Indeed, susceptibility to the trans inclusion problem is a general feature of social-position accounts. If an individual’s gender is wholly determined by facts about other individuals, some individuals who identify as women may not belong to Gender 1 and some individuals who identify as men may not belong to Gender 2.

15 Bettcher (2009), for example, articulates a conception of gender as “existential self-identity.” See Barnes (2022, 838) for a list of scholars and advocacy organizations which have embraced the view that gender identification is the sole determinant of gender.

16 For another recent proposal, see Ashley (2023).

17 See Jenkins (2018) for further discussion.

18 Thus, though the term “woman” appears on the right-hand side of the biconditional in Jenkins’s definition, her account is not circular because that term can be paraphrased away. Nor, as I interpret it, is Jenkins’s account of gender identity trans-exclusionary, as Bogardus (2020, 62) argues: “the trans women who are excluded from Haslanger’s definition can’t sensibly take norms for women-in-Haslanger’s-sense to be norms for *them*, and so they can’t sensibly be women on Jenkins’ view” (emphasis in original). As I understand Jenkins’s map metaphor, having a given internal map or taking a given set of norms to be relevant to oneself is not a matter of forming a belief, sensible or otherwise, to the effect that one falls into the category of individuals for which the norms are intended. It is something closer to a cognitive or experiential disposition. For example, Jenkins writes: “Consider a woman who feels that having visible body hair on her legs is unattractive, embarrassing, and unacceptable. In a visceral way, having hairy legs feels wrong for her. This feeling—this instinctive sense of how her body ‘ought to be’—is part of her gender identity.” (2016, 411).

19 Barnes also articulates what she calls “more general exclusion worries” for identity accounts of gender (2020, 711). Her basic idea is that identity accounts are faced with a dilemma: if they make the conditions for having a given gender identity too stringent, they will exclude some gender non-conforming women; if they make the conditions for having a given gender identity too lax, they will include some non-women (genderqueer people, for example, or drag performers). I agree with Barnes that any adequate account of gender identity will need to walk a fine line between overgeneration and undergeneration, but—unlike the disability inclusion problem, which seems to me inescapable—these more general exclusion worries strike me as desiderata that an adequate account of gender identity will have to address rather than as objections to the idea of gender as identity.

20 The gender inclusion problem also arises for a number of proposals which resist characterization as either social-position accounts or identity accounts. Stoljar (1995, 2011), for example, defends a resemblance-nominalist position, according to which to be a woman is to possess a sufficient number of the features in our concept *woman*, which is a cluster concept with elements including gender identity, social position, sex, and lived experience. Stoljar resists the idea that there are necessary or sufficient conditions for womanhood. This means that it is difficult for her view to escape the trans inclusion problem, since some trans women are intuitively women solely in virtue of their gender identities. Along similar lines, McKittrick (2015) defends the position that an individual’s gender is partially grounded in their internal dispositions to behave in certain masculine or feminine ways. This exposes her account to the disability inclusion problem. And “gender-critical” approaches (e.g. Lawford-Smith 2022) face the trans inclusion problem by design in so far as they hold that the explanatory and emancipatory goals of a feminist metaphysics of gender are best served by adopting a theory of gender categories which ties them extensionally to sex and thereby excludes trans women from Gender 1. Lawford-Smith, for example, writes: “The feminists who think gender is a performance or an identity appear to believe that there is no loss to women in shifting to an alternative understanding of gender, and therefore a new constituency for feminism (and in many places, redrawn boundaries of the group eligible for protection as female/woman under the law and in social policy). The aim of this chapter is to show that they are wrong.” (2022, 94). Of course, by their own lights, gender-critical feminists do not face the trans inclusion problem, since they reject the assumption that it is desirable to develop a trans-inclusive theory of gender. It is beyond the scope of my discussion here to provide a sustained defense of the project of trans-inclusive feminist accounts of gender. Dialectically, my aim is to enter into a discussion internal to trans-inclusive feminism and show how to pursue the project of constructing a metaphysical account of gender in a way that solves inclusion problems.

21 Another approach to the gender inclusion problem is to hold that trans or disability exclusionary theories of gender categories reflect a problem with the world, not a problem for the feminist metaphysician. This approach has been influentially defended by Ásta (2018), who distinguishes sharply between empirical questions about the metaphysics of gender and normative questions about whether the way we construct gender is justified: “we could answer the question *what is a woman?* empirically: in this and that context, the base property for the conferral of the status is sex assignment, or the presence of a vagina, or self-identification. But that doesn’t *settle* the matter. Is the base property in question justified?” (Ásta 2018, 91; emphasis in original). As I understand theorists like Ásta, they adopt a methodology different from the one I have articulated above in that they reject including trans and disability inclusivity among the emancipatory desiderata on a feminist theory of gender categories. According to this methodology, the feminist metaphysician can set aside normative questions about how social reality ought to be constructed when giving a theory of gender category membership. Worries about trans and disability inclusivity come into play when we think about what kinds of gender categories it would be *best* to socially construct, not when we

think about what kinds of gender categories we have actually constructed. In one sense, then, theorists like Ásta sidestep the gender inclusion problem.

In another sense, however, they do not. For as soon as we ask which gender categories it would be best to socially construct, we are faced with the problem that existing proposals are subject to either the trans inclusion problem or the disability inclusion problem. A theory that solves both problems at once, like the one I offer, should therefore be of interest even to theorists like Ásta as a description of an inclusive practice of gendering that can serve as a target for ameliorative interventions. Thanks to an anonymous referee for pressing me to clarify this point.

22 Though it is a structured account, (Inclusive Genders) gives *identity priority* in the sense that whether an individual is feminized or masculinized determines their gender only if they lack a gender identity. An alternative would be to give social position priority, as in:

(Inclusive Genders*): An individual belongs to Gender 1 just in case they are feminized in a suitably broad range of contexts, unless they form a gender identity which conflicts with belonging to Gender 1, in which case their gender is determined by this identity. The same is true, *mutatis mutandis*, for Gender 2.

(Inclusive Genders) and (Inclusive Genders*) issue the same verdicts about most cases, and I do not wish to claim that one is obviously superior to the other—both are structured theories of gender with the resources to solve the gender inclusion problem. Nevertheless, I think there are two considerations which favor (Inclusive Genders) over (Inclusive Genders*). First, I think it is plausible that, in cases where an individual has a gender identity, this identity grounds facts about their gender. For this reason, I think (Inclusive Genders) is more suggestive of the underlying metaphysical structure than (Inclusive Genders*)—though both make claims about necessary and sufficient conditions rather than grounding claims. Second, if it is possible for an individual to have a gender identity but not a social position, the two accounts make different predictions. Imagine, for example, that a person who is born and grows up alone on a desert island discovers a stock of lifestyle magazines and comes to identify with the women she finds depicted in them. (Inclusive Genders) correctly predicts that this individual belongs to Gender 1, while (Inclusive Genders*) predicts that she belongs to neither gender because she is not feminized in a suitably broad range of contexts and her gender identity does not conflict with belonging to Gender 1.

23 Though the term “Gender 1” appears on the right-hand side of the biconditional in (Inclusive Genders), the account is not circular because that occurrence of “Gender 1” can be paraphrased away. While my case for (Inclusive Genders) does not turn on any commitments concerning genders beyond Gender 1 and Gender 2, many may regard it as an advantage of the structured account that it makes room for further gender categories. In fact, it does so in two ways. First, it makes room for further gender categories that are like Gender 1 and Gender 2 in that their existence is grounded in an underlying social practice of positioning individuals on the basis of what is observed or imagined about them. Identifying with such a category would then be a necessary and sufficient condition for belonging to it for those with a gender identity, while being socially positioned in it would be the criterion for membership for individuals lacking a gender identity. This may be the correct metaphysical story about cultural contexts which traditionally recognize a third gender, for example the Samoan *fa’afafine* and the Indian *hijra*. It is also plausibly what George and Briggs (2019, 27) have in mind when they write that “If some intersex conditions (or families of such conditions) are likewise understood as sexes in a given society, then that society might associate them with distinct primordial gender categories.”

Second, it makes room for further gender categories membership in which is grounded wholly in facts about identity. This is one way to understand the category *nonbinary*—that is, as the gender category to which an individual belongs just in case they have a gender identity but do not identify with either Gender 1 or Gender 2. Something like this might also be what Dembroff (2020, 16) has in mind when they analyze *genderqueer* as an identity based on “a felt or desired gender categorization that conflicts with the binary axis” that “collectively destabilizes this axis,” or what George and Briggs (2019) have in mind when they write that “The *genderfluid* identity is not plausibly analyzed, or even idealized, as being rooted in any primordial gender category. Instead, what makes it a gender is the way that it is understood within a system of existing gender categories. If somebody sometimes conceives of and presents themselves as a woman, or like a woman, and other times conceives of and presents themselves as a man, or like a man, this makes them an exemplar of (one sort of) genderfluidity” (28). As I understand them, neither Dembroff’s category *genderqueer* nor George and Briggs’s category *genderfluid* requires the existence of a dominant social practice of recognizing *genderqueer* or

genderfluid as gender categories on a par with Gender 1 and Gender 2. Thanks to an anonymous referee for inviting me to elaborate on this point.

24 A fruitful source for further examples of disjunctive social properties is the anthropological literature on kinship relations. As Marshall Sahlins (2013, 2), Sahlins details a number of examples of kinship properties in cultures around the world that appear to be disjunctively grounded either in biological relationships of heredity or in purely social facts (e.g., facts about ceremonies, cohabitation, shared experiences, and so forth).

25 Thanks to an anonymous referee for suggesting this response.

26 Along similar lines, (Inclusive Genders) has the resources to explain what is wrong with treating nonbinary people as though they belong to Gender 1 or Gender 2. Since identity takes priority in determining gender category membership, nonbinary people belong to neither Gender 1 nor Gender 2. So treating a nonbinary person as though they belong to Gender 1 or Gender 2 is making a mistake, and this is so even if that person is feminized or masculinized by others across a range of contexts. One helpful way of understanding what (Inclusive Genders) predicts in this domain is to appeal to Dembroff's (2018) idea of *ontological oppression*, one variety of which occurs when "The structures and practices within a social context . . . unjustly fail to recognize or construct certain kinds" (25). If (Inclusive Genders) is correct, then treating nonbinary people as though they belong to Gender 1 or Gender 2 is a form of ontological oppression in so far as it involves failing to recognize that they belong to neither social kind. Thanks to an anonymous referee for discussion on this point.

27 This interpretation is suggested by Bettcher's remark that "the expression 'non-trans woman' operates in the way that 'woman' used to operate" (2013, 241).

28 As discussed in n. 23 above, some cultures have constructed social reality in this way—contemporary examples include the Samoan *fa'afafine* and perhaps also the Indian *hijra*.

29 At the same time, Barnes's metaphysical proposal is not in competition with (Inclusive Genders); indeed, the two could form complementary parts of a total theory of gender and gender categories according to which the existence of gender is grounded in the way Barnes indicates and the membership conditions of particular gender categories are given by (Inclusive Genders). So Barnes's proposal is consistent with a solution to the gender inclusion problem. But in this total theory of gender and gender categories, it would be (Inclusive Genders), rather than the Barnes-style social-position account, which would solve the gender inclusion problem.

30 In what follows, I focus on the meanings of the gender terms "woman" and "man." Of course, there are many other linguistic expressions which implicate gender. For recent discussion, see Hernandez and Crowley (2024), González Vázquez et al. (2024), and Kirk-Giannini and Glanzberg (2024).

31 A contextualist theory of this kind is endorsed by Barnes (2020) and discussed (though not endorsed) by Saul (2012).

32 Barnes (2020) endorses a theory of this kind.

33 The corresponding inclusion problem for gendered language pertaining to disabled individuals can be formulated similarly. Note that I have intentionally chosen the open-ended locution "an account of our linguistic practice . . . fails to secure" in order to abstract away from differences between semantic and pragmatic explanations and leave open both approaches which, like Barnes's, seek to describe what is true of our actual linguistic practice and approaches which, like Jenkins's, seek to describe what is true of the linguistic practice we have most reason to adopt.

As was the case with the gender inclusion problem, the gendered language inclusion problem is a problem only from a certain theoretical and methodological perspective. Those who reject the idea that the emancipatory goals of feminism recommend ensuring that our linguistic practices are trans and disability inclusive will not feel the force of the problem as an objection to certain proposals about how we do or ought to use gendered language. It is beyond the scope of my discussion here to argue for the theoretical and methodological perspective which gives rise to the gendered language inclusion problem—it will have to suffice simply to note that many feminists occupy that perspective, and that the solution suggested by (Inclusive Genders) will therefore be of interest to them.

34 Indeed, there is a lively debate about whether contextualist accounts of gendered language can be adequately trans-inclusive. See, e.g., Saul (2012), Diaz-Leon (2016), Barnes (2020), Ichikawa (2020), Laskowski (2020), Zeman (2020), and Chen (2021).

35 To be clear, I do not offer (Inclusive Extensions) as a description of our actual linguistic practice.

36 Identifying (Inclusive Extensions) as a linguistic practice which should be adopted in light of the feminist's emancipatory goals also gives us a straightforward explanation of why it is inappropriate to describe a nonbinary person as a man or a woman: because they belong to neither Gender 1 nor Gender 2, they ought to fall into the extension of neither "woman" nor "man." Thanks to an anonymous referee for discussion on this point.

References

- Antony, Lousie. 2020. Feminism without metaphysics or a deflationary account of gender. *Erkenntnis* 85: 529–49.
- Ashley, Florence. 2023. What is it like to have a gender identity? *Mind*. Online First.
- Ásta. 2011. The metaphysics of sex and gender. In C. Witt (ed.) *Feminist Metaphysics*, ed. C. Witt. New York: Springer.
- Ásta. 2013. The social construction of human kinds. *Hypatia* 28: 716–32.
- Ásta. 2018. *Categories we live by*. Oxford: Oxford University Press.
- Barnes, Elizabeth. 2020. Gender and gender terms. *Noûs* 54: 704–30.
- Barnes, Elizabeth. 2022. Gender without gender identity: The case of cognitive disability. *Mind* 131: 836–62.
- Beauvoir, Simone de. 1953. *The second sex*. London: Jonathan Cape.
- Bernstein, Sara. 2020. The metaphysics of intersectionality. *Philosophical Studies* 177: 321–35.
- Bettcher, Talia Mae. 2009. Trans identities and first-person authority. In L. Shrage (ed.), "You've changed": *Sex reassignment and personal identity*, ed. L. Shrage. Oxford: Oxford University Press.
- Bettcher, Talia Mae. 2013. Trans women and the meaning of "woman." In *The philosophy of sex: Contemporary readings*, 6th ed., ed. N. Power, R. Halwani, and A. Soble. Lanham, MD: Rowman & Littlefield.
- Bogardus, Tomas. 2020. Some internal problems with revisionary gender concepts. *Philosophia* 48: 55–75.
- Butler, Judith. 1990. *Gender trouble*. New York: Routledge.
- Butler, Judith. 1993. *Bodies that matter*. New York: Routledge.
- Byrne, Alex. 2020. Are women adult human females? *Philosophical Studies* 177: 3783–3803.
- Byrne, Alex. 2021. Gender muddle: Reply to Dembroff. *Journal of Controversial Ideas* 1 (5): 1–24.
- Byrne, Alex. 2022. The female of the species: Reply to Heartsilver. *Journal of Controversial Ideas* 2 (11): 1–22.
- Chen, Hsiang-Yun. 2021. Contextualism and the semantics of "woman." *Ergo* 8: 577–97.
- Dembroff, Robin. 2018. Real talk on the metaphysics of gender. *Philosophical Topics* 46: 21–50.
- Dembroff, Robin. 2020. Beyond binary: Genderqueer as critical gender kind. *Philosophers' Imprint* 20 (9): 1–23.
- Dembroff, Robin. 2021. Escaping the natural attitude about gender. *Philosophical Studies* 178: 983–1003.
- Diaz-Leon, Esa. 2016. *Woman* as a politically significant term: A solution to the puzzle. *Hypatia* 31: 245–58.
- George, B. R., and R. A. Briggs. 2019. Science fiction double feature: Trans liberation on Twin Earth. Manuscript.
- González Vázquez, Iz, Anna Klieber, and Martina Rosola. 2024. Beyond pronouns: Gender visibility and neutrality across languages. In *The Oxford handbook of applied philosophy of language*, ed. L. Anderson and E. Lepore. Oxford: Oxford University Press.
- Haslanger, Sally. 2000. Gender and race: (What) are they? (What) do we want them to be? *Noûs* 34 (1): 31–55.
- Heartsilver, Maggie. 2021. Deflating Byrne's "Are women adult human females?" *Journal of Controversial Ideas* 1 (9): 1–16.
- Hernandez, E. M., and Archie Crowley. 2024. How to do things with gendered words. In L. Anderson and E. Lepore (eds.) *The Oxford handbook of applied philosophy of language*, ed. L. Anderson and E. Lepore. Oxford: Oxford University Press.
- Ichikawa, Jonathan J. 2020. Contextual injustice. *Kennedy Institute of Ethics Journal* 30: 1–30.
- Jenkins, Katharine. 2016. Amelioration and inclusion: Gender identity and the concept of *woman*. *Ethics* 126: 394–421.
- Jenkins, Katharine. 2018. Toward an account of gender identity. *Ergo* 5: 713–44.
- Jenkins, Katharine. 2023. *Ontology and oppression: Race, gender, and social reality*. Oxford: Oxford University Press.

- Kirk-Giannini, Cameron Domenico, and Michael Glanzberg. 2024. Pronouns and gender. In *The Oxford handbook of applied philosophy of language*, ed. L. Anderson and E. Lepore. Oxford: Oxford University Press.
- Laskowski, Nick G. 2020. Moral constraints on gender concepts. *Ethical Theory and Moral Practice* 23: 39–51.
- Lawford-Smith, Holly. 2022. *Gender-critical feminism*. Oxford: Oxford University Press.
- McKittrick, Jennifer. 2015. A dispositional account of gender. *Philosophical Studies* 172: 2575–89.
- Mikkola, Mari. 2022. *Feminist perspectives on sex and gender*. Stanford Encyclopedia of Philosophy. <https://plato.stanford.edu/entries/feminism-gender/>.
- Sahlins, Marshall. 2013. *What kinship is—and is not*. Chicago: University of Chicago Press.
- Saul, Jennifer. 2012. Politically significant terms and the philosophy of language: Methodological issues. In *Out from the shadows: Analytical feminist contributions to traditional philosophy*, ed. S. L. Crasnow and A. M. Superson. Oxford: Oxford University Press.
- Stoljar, Natalie. 1995. Essence, identity, and the concept of woman. *Philosophical Topics* 23: 261–93.
- Stoljar, Natalie. 2011. Different women: Gender and the realism–nominalism debate. In C. Witt (ed.) *Feminist Metaphysics*, ed. C. Witt. New York: Springer.
- Witt, Charlotte. 2011. *The metaphysics of gender*. Oxford: Oxford University Press.
- Zeman, Dan. 2020. Subject-contextualism and the meaning of gender terms. *Journal of Social Ontology* 6: 69–83.

Cameron Domenico Kirk-Giannini is an assistant professor of philosophy at Rutgers University–Newark working on a range of issues in social philosophy, philosophy of language, and philosophy of artificial intelligence.