

Normative Metaphysics

Dee Payton

Draft June 2024

Abstract. Many projects in feminist metaphysics appear to background something like a principle of “ought implies is.” This is a principle which takes us from claims about what *should be* the case given the demands of justice, to claims about what *is* the case, ontologically speaking. In this paper, I’m interested to know how that principle is best interpreted within the context of these projects and what can be said in its defense. To this end, I explore three interpretations of the principle: (i) that normative commitments generate epistemic constraints on theories in feminist metaphysics; (ii) that gender terms like ‘woman’ turn out to be normative, in much the same way as terms like ‘good’ and ‘just’; and (iii) that the normative commitments of these projects somehow directly deliver their ontological posits. I call these the epistemic interpretation, the metaethical interpretation, and the ontological interpretation of the “ought implies is” principle, respectively.

1 Introduction: Normative Authority in Feminist Metaphysics

The idea that our values can and often do play a role in philosophical theorizing is not new. Consider, for example, the role of theoretical virtues like simplicity. There is some question, of course, about what simplicity involves and what its demands require. But traditionally, many philosophers have endorsed some version of the claim that, given two theories which explain the same phenomenon, if the first accomplishes this task in a way which is simpler than the second, we should prefer the first. In at least this general sense, our theories should be simple.

There is also the familiar notion that, in conducting theoretical investigations, we're interested not only in uncovering truths, but in revealing explanatorily significant truths. There is a sense in which truth of a general sort is easy to come by, but the notion of explanatory significance embeds the idea that some truths are more important than others. Of course, philosophers disagree over what makes for explanatory significance. But most ways of answering this question have it that there are special truths out there, and there is a sense in which we should theorize in terms of them.

Now consider the idea that we ought to theorize in the service of justice. This is also a way that values can be incorporated into theory. And presumably, there are a great many roles that a concern for justice could play in guiding our theorizing. Here, I'm interested in one particular way that our theories can be informed by what justice demands, and this is in the move from claims about what should be the case, to claims about what is the case. Moves like this are common in feminist metaphysics. To see this, consider three prominent projects in the metaphysics of gender:

Sally Haslanger has argued for a social position account of gender.¹ Very generally, on her view, someone is a woman just in case they are subordinated in certain ways because they are taken to be biologically female, and someone is a man just in case they are privileged in certain ways because they are taken to be biologically male. Haslanger argues that these properties are *gender* properties in virtue of the role they play in explaining the (continued) existence of certain

¹ Sally Haslanger, *Resisting Reality* (Oxford: Oxford University Press 2012).

hierarchical social structures. It is also in virtue of playing this explanatory role that these properties are among the meanings of our gender terms ‘woman’ and ‘man’. Of interest here is how Haslanger singles out the particular explanatory role played by these properties. For, one might ask: why care about *this* particular explanatory role, rather than *that* one? Haslanger’s answer is, in essence, that this is the explanatory role feminists *should* care about, given our interest in bringing about a just world. In other words, this explanatory role matters given Haslanger’s background commitments to justice.

In some contrast to Haslanger’s approach, Esa Díaz-León has argued for a deflationary approach to the metaphysics of gender.² On their view, questions about the metaphysical reality of gender are best treated through a combination of conceptual analysis and empirical investigation. So for example: where the question is “what is it to be a woman?” we proceed here by (i) determining the meaning of ‘woman’, and then (ii) looking out into the world to see whether ‘woman’, so defined, refers to anything. Importantly, this does not mean that we can define our gender terms and concepts however we want, and so populate our gender ontology however we like. There are constraints here, namely: in determining the meanings of our gender terms, we’re interested mainly in what these terms *should* mean, given our theoretical purposes, which many times are *normative* purposes. Which ways of assigning meanings to gender terms and concepts would be the most explanatorily useful, given our aims? Which moral and political considerations are the most relevant when it comes to assigning those meanings? Those are the questions Díaz-León argues that we have been and should continue to be focusing on in the philosophy of gender, and importantly, those questions position our normative commitments in an authoritative role.

Lastly, consider Katharine Jenkins’ twin target account of gender.³ Jenkins’ account is largely motivated by a trans-inclusive desideratum on theories of gender, a desideratum which requires

² Esa Díaz-León, “*Woman as a Politically Significant Term: A Solution to the Puzzle*,” *Hypatia*, 31, 2 (2016) :245-258; Esa Díaz-León, “On Haslanger’s Meta-Metaphysics: Social Structures and Metaphysical Deflationism,” *Disputatio*, 10, 50 (2018): 201-216; Esa Díaz-León, “Substantive metaphysical debates about gender and race: Verbal disputes and metaphysical deflationism,” *Journal of Social Philosophy*, 53, 4, (2021): 562.

³ Katharine Jenkins, “Amelioration and Inclusion: Gender Identity and the Concept of Woman,” *Ethics*, 126, 2 (2016): 394-421.

a feminist analysis of gender to be one which respects the identifications of trans people by including them in the gender categories with which they identify (and not including them within any categories with which they do not identify). Of interest here is the fact that Jenkins takes this desideratum to follow directly from considerations about the oppression and marginalization of trans people, in view of considerations of justice. In other words, Jenkins' reasoning here is as follows: given that trans people are oppressed and marginalized in multiple respects, it follows that any feminist account of gender should satisfy this desideratum. Jenkins argues that their dual account of womanhood satisfies this trans-inclusive desideratum for theories of gender, and this is a strong reason to endorse their view.

Now of course, the projects described above differ in a number of significant respects. But they share a common methodology: namely, the central ontological commitments of each project somehow proceed from its underlying normative commitments. In this essay, I'm interested in how that move works.

Given the diversity of views in the literature, it's likely that even after narrowing the scope of our attention here, there are still several ways to understand the relationship between the normative and ontological commitments of these projects. Here, I'll examine one particularly strict way of understanding this relationship, in terms of a move from claims about what ought to be the case, to claims about what is the case, the sort of move made legitimate by something like a principle of "ought implies is." Now, for many projects in the literature, this will likely be a somewhat oversimplified and overpowered reading of this relationship. But a discussion of this principle in its most basic form promises to be instructive. For of course, if a basic reading of the principle fails, it will be very good to have a sense for both how and why, and what the consequences that failure might be for how we understand different versions of this methodology in the existing literature.

Here, I will proceed under the assumption that a version of this principle is working in the background of projects in feminist metaphysics like those mentioned above. This is a principle which allows us to move from claims about what should be the case (given the demands of justice) to claims about what is the case (ontologically speaking). I'm interested to know how a principle like that is best interpreted and what might be said in its defense. To this end, I won't

limit my discussion to any one particular project. Instead, I'll explore this general methodological move within the context of questions about gender, which are among the central questions in this literature. These are questions like "what is it to be a woman?", "what is it to be a man?" and "what is it to be nonbinary?". Here, I explore how this principle of "ought implies is" might operate in the background of philosophical inquiries centered on questions like those.

I've decided to hold this discussion within the context of feminist metaphysics for a few reasons. First and foremost, this general methodological move is common in that literature. Second, despite the prominence of this methodology in feminist metaphysics, this move hasn't received much direct attention; if that's right, then there is a significant and foundational gap here, which I hope my discussion in this paper can do some work to fill. Third, I am a feminist metaphysician and so I have something like a personal stake in the outcome of this exploratory project (just laying my cards on the table here). And it is an exploratory project. My general aim in this paper is not to argue for or against the methodological move I investigate here, but simply to explore possible ways of cashing out how it could work.

In the sections below, I explore three interpretations of the "ought implies is" principle: (i) that normative commitments generate epistemic constraints on theories in feminist metaphysics; (ii) that gender terms like 'woman' turn out to be normative, in much the same way as terms like 'good' and 'just'; and (iii) that the normative commitments of these projects somehow directly deliver their ontological posits. I call these the epistemic interpretation, the metaethical interpretation, and the ontological interpretation of the "ought implies is" principle, respectively.

As a bit of insight into the outcome of this expedition: ultimately, I conclude none of these interpretations of the principle succeeds. Success here requires two things: first, the proposed interpretation must "*fit*" the context of inquiry. In other words, it must accomplish the work that feminist metaphysicians need it to perform. And second, if a particular interpretation "fits," it must also be *defensible*, in the sense that it holds up under scrutiny. Admittedly, this second success criterion is loosely defined, but as we'll see, several interpretations of "ought implies is"

fail to satisfy even the most generous readings of this criterion, insofar as they leave the connection between our normative and ontological commitments wholly opaque.

Some readers may benefit from seeing my exploratory project here phrased in terms of a positive argument. And so, the central claims of this paper can be put as follows:

- (i) First, there is something like a principle of “ought implies is” working in the background of many projects in feminist metaphysics.
- (ii) Second, this principle can be interpreted in at least three ways (these are the epistemic, metaethical, and ontological interpretations of the principle).
- (iii) Third, each of the proposed interpretations fails.

Here's a brief sketch of where we're headed next: in Sections II and III, I'll argue that the question I'm asking here is an important one, and it is an open one—that is, this question matters and it hasn't been answered yet. Then in Section IV, I'll discuss the three interpretations of “ought implies is” introduced above. Finally, in the last section of this paper, I'll briefly turn to a more general discussion regarding the degree to which the theoretical aims of projects in feminist metaphysics align with those of projects in traditional analytic metaphysics. I argue that one of the most promising ways of closing the gap between our normative and ontological commitments in feminist metaphysics may require a significant departure from the truth-seeking aims of more traditional metaphysics, and this raises a question about the methodological compatibility of the two disciplines.

2 The Problem of Missing Value

In this section and the next, my goal is to both isolate and motivate the question at the heart of this paper, regarding the relationship between normative and ontological commitments in feminist metaphysics. My claim is that there really is an important question here which hasn't yet been answered by the existing literature.

I want to make this discussion accessible to a general audience, but with this argument I also have two particular groups of philosophers in mind: those with a background in what might be called “mainstream” metaphysics, who may be familiar with recent work on the role of value in

metaphysical theorizing; and those with a background in feminist philosophy and critical theory, who may be familiar with work in feminist epistemology and philosophy of science on value-laden approaches to theory selection. The first part of my argument addresses work in “mainstream” metaphysics; the second part of my argument addresses work in feminist epistemology and philosophy of science.

Shamik Dasgupta has recently argued that realist projects in metaphysics struggle with what he calls *the problem of missing value*.⁴ Here, generally, is the problem: metaphysical realist views posit what Dasgupta calls “elite” properties. A property is elite just in case it is theory-guiding, in a certain sense. For example: the elite properties at work in Sider’s structural realist view are natural properties; natural properties are theory-guiding, and so elite, insofar as we should theorize in terms of them, according to Sider.⁵ In this way, “eliteness is by definition a value-theoretic phenomenon: to be elite is to be something our theorizing *should* center around.” Dasgupta uses this observation about the role of elite properties in realist projects to distinguish two claims which he takes to be central to these projects. These are “a purely *metaphysical* claim that posits some objective property of ‘being a law,’ or relation of grounding, or whatever; and a *value-theoretic* claim to the effect that the metaphysical posit functions as an objective standard against which our theorizing is to be assessed.”⁶ However, the problem here is that realists do not provide a defense of the value-theoretic claim, but simply assume that it follows, given the truth of the metaphysical claim. In the context of Sider’s project, for example, this looks like assuming that natural properties should be theory guiding (the value-theoretic claim), given that there are natural properties (the metaphysical claim). But why should we think that the value theoretic claim follows from the metaphysical claim here? This is the question at the heart of Dasgupta’s paper, and he argues that realist projects like Sider’s don’t contain the resources to answer it.

The question I’m exploring in the present paper is related to, but importantly distinct from, the question Dasgupta discusses. In essence: Dasgupta is interested in *why* we should think that

⁴ Shamik Dasgupta, “Realism and the Absence of Value,” *Philosophical Review*, 127, 3 (2018): 279-322.

⁵ Theodore Sider, *Writing the Book of the World* (Oxford: Oxford University Press 2011).

⁶ See fn. 4: Dasgupta (2018), p. 283.

these value-theoretic claims are true, whereas I am interested in *how* value-theoretic claims guide our theorizing, given that these claims hold. In this way, our projects naturally accompany one another, a relationship Dasgupta also acknowledges, as evidenced by the fact that, in framing his discussion of the problem of missing value, he writes “[to] be theory-guiding is to play a normative role in theorizing, but what exactly is this role? This is a good question and not one that I can hope to answer here.”⁷ In brief: I agree that this is a good question, and exploring answers to it is the project of the present paper.⁸

3 Feminist Methodology and Value-Laden Inquiry

I hope that the foregoing section has done something to clarify and motivate my central line of inquiry here. However, I anticipate that some readers will remain skeptical that this question is an interesting one to ask, specifically within the context of feminist metaphysics. For, one might be persuaded by Dasgupta’s discussion about the role of value-theoretic claims in “mainstream” metaphysics, but think that once we shift our focus to projects in feminist philosophy, this question becomes uninteresting, insofar as the aims of theoretical inquiry in this context are importantly different.

To appreciate this angle, start with the observation that, to the extent that feminist metaphysicians have dealt with the question of how normativity influences their ontological commitments, it has largely been with reference to the value-laden methodological frameworks articulated in feminist epistemology and philosophy of science. This value-laden approach to

⁷ See fn. 4: Dasgupta (2018), p. 290.

⁸ Answering it within the context of feminist metaphysics, that is. For, this discussion of Dasgupta’s project suggests that there is a much more general question to be asked here about the role of value theoretic claims in “mainstream” metaphysics more broadly. I will not attempt to address this more general version of the question however, because I have the sense that both the nature and role of value-theoretic claims in projects like Sider’s differ significantly from the nature and role of these claims as they are posited by projects in feminist metaphysics (more on this in Section V). I think these differences make it difficult hold a discussion which does justice to both sorts of projects at the same time (at least, given the constraints on space within a single essay), and so I focus exclusively on projects in feminist metaphysics here, leaving it open that some of what I say may generalize.

theory selection can take a number of different forms, but I will focus on Elizabeth Anderson's view here because her work has consistently been cited within feminist metaphysics.⁹

Anderson has argued for what she calls *value-laden* approaches to theory selection. This is a methodology according to which, roughly, empirical truths taken alone underdetermine theory choice, and context-dependent values must be added to the mix in order to justify the choice of a given theory over available alternatives. This value-laden methodology seeks to answer our question regarding the relationship between normative and ontological commitments with a distinction between truth and justification in the context of theory choice. Generally: in investigating the nature of e.g., gender, we want true, *justified* theories; metaphysical truths alone can offer us a range of candidate true theories, but won't uniquely privilege any one candidate. Our contextual values—our underlying normative commitments—do the work of justifying one theory above the rest. So, where our question is one of how normative truths might deliver ontological ones, according to this approach, the answer is that normative truths *justify* the theories according to which those ontological claims hold.

As part of her case for this approach, Anderson directly addresses arguments against value-laden methodology. In this discussion, she argues that debates over “ought implies is” turn out to be a red herring, and that the central point of tension here manifests in a disagreement about the aims of theoretical inquiry. For, proponents of value-laden methodology will argue that truth is *not* the sole aim of theoretical inquiry, and that a theory can be justified by factors in addition to truth.¹⁰

There is much to be said in favor of Anderson's approach here. But I disagree that debates about “ought implies is” turn out to be a red herring. This is because, even if Anderson's

⁹ See Elizabeth Anderson, “Knowledge, Human Interests, and Objectivity in Feminist Epistemology,” *Philosophical Topics*, 23, 2 (1995a): 27-58; Elizabeth Anderson, “Feminist Epistemology: An Interpretation and Defense,” *Hypatia*, 10, 3, (1995b): 50-84. Elizabeth Anderson, “Uses of Value Judgments in Science: A General Argument, with Lessons from a Case Study of Feminist Research on Divorce,” *Hypatia*, 19, 1 (2004): 1-21.

¹⁰ To the extent that projects in feminist metaphysics are best understood as making use of Anderson's methodology here, this claim about the aims of theoretical inquiry foreshadows the conclusions I'll discuss in Section V, regarding the divide between feminist metaphysics and more traditional metaphysics.

argument against value-neutral methodology is successful, her positive framework still takes us to a version of “ought implies is.” It’s just that, on her view, this relationship is phrased in slightly different terms. I’ll use a generalized example to begin my argument here—

Say that we arrive at a theory of gender via Anderson’s methodology for theory choice. Importantly, according to Anderson, values enter into our theorizing at a number of different places. For instance: in our decision to theorize about gender in the first place; in our formulation of the questions which frame our inquiry; in our approach to collecting and analyzing evidence and relating it back to our starting hypothesis; in formulating the content and consequences of the resulting theory; and in evaluating that theory with respect to available alternatives. Let’s say we take all of that into account, and via this process, arrive at a particular theory of gender, Φ .

But in what sense do these values make Φ our theory of gender? Note that this question concerns what it is for value claims (and in particular, claims about what should be the case) to *justify* theories, if we aren’t supposed to understand that justification in terms of either truth or a further normative claim (e.g., something like Φ should be true or Φ should be our theory of gender). That is, on Anderson’s view: for a set of values V to justify a theory Φ is not for V to support the *truth* of Φ , nor for V to support claims about how Φ *should* be employed, but something else.

Some of Anderson’s later arguments suggest that for Φ to be justified by V is for Φ to be the most *useful* theory given V , insofar as Φ promises to be e.g., empirically fruitful and pragmatically helpful in certain respects that matter to us, given our values. On this approach, particular facts about gender are true because they hold according to our most useful theory of gender.

However, the way I see it, this approach only works to bridge the gap from normative claims to ontological ones *if we are pragmatists about the nature of (social) truth*. That is, this approach succeeds only if we are committed to a view of truth according to which for particular facts about gender to be true is for those claims to have some pragmatic utility. For, if we are not pragmatists about truth in this way, then the question I’ve been pressing here arises yet again: for, how does it follow from Φ ’s being the most *useful* theory of gender that e.g., a particular property *is* the

property of being a woman (unless we understand that latter ontological claim in pragmatic terms as well)?

I think it would be best if feminist metaphysicians didn't *have* to be pragmatists about the nature of truth in order to have our normative commitments play the authoritative role we'd like them to in these projects.¹¹ It is also my sense that some of the feminist philosophers quoted in the introduction (and Haslanger stands out here!) have projects with elements which appear fundamentally incompatible with this sort of thoroughgoing pragmatism.¹²

With that said, this sort of pragmatism may turn out to be our only option, when it comes to the task of bridging the gap between the normative and ontological commitments at the heart of these projects in feminist metaphysics. I will return to this point later on, as my sense is that ultimately, questions about the viability of pragmatism in this context direct us toward a much larger methodological question about the aims of inquiry in feminist metaphysics, and whether/how those aims might differ from those of projects in more traditional analytic metaphysics (Section V). Note, however, that if we bracket this discussion of pragmatism, Anderson's methodology straightforwardly faces the challenge posed by "ought implies is" for the reasons described above, and so we're right back where we started.

There is one last route to avoiding the challenge posed by "ought implies is" that I'd like to explore here, a strategy which arises from the observation that Anderson's view seems to be premised on a deflationary metaphysical worldview.¹³ A commitment to the parity of metaphysical distinctions is a hallmark of deflationist worldviews. Whereas on a realist picture, the world comes pre-built with a deep metaphysical structure, one which our classificatory practices can do a better-or-worse job of tracking, deflationists are committed to the view that

¹¹ John Capps, "The Pragmatic Theory of Truth," in Edward N. Zalta & Uri Nodelman (eds.) *The Stanford Encyclopedia of Philosophy* (Summer 2023).

¹² I will add to this that I personally find it deeply unappealing to say that, e.g., the truth of [women in Afghanistan are oppressed by the Taliban] amounts to that claim having a certain level of pragmatic utility, given our values. If anything, that claim has pragmatic utility because it is true, and it is true because that's the way the world is, not the other way around! Relatedly, see Barnes' (2017) "direction of explanation" objection to deflationist theories of gender: Elizabeth Barnes "Realism and Social Structure," *Philosophical Studies* 174, 10 (2017): 2417-2433.

¹³ See fn. 9: Anderson (1995a), p. 45.

we bring that structure to the world with our values, and our linguistic and conceptual representations. And importantly, there is no deep structure to reality outside of this.¹⁴

Now, it might be tempting to think that metaphysical deflationists have an easier time than realists do when it comes to handling this question regarding “ought implies is.” For, if we are deflationists about e.g., gender, then our ontological claims about gender have a different status than those of the realist. Whereas a realist might appeal to deep metaphysical structures in arguing for their view of gender,¹⁵ when a deflationist says that a property F is the property of being a woman, all they are saying is that F is a (proper) meaning of the gender term ‘woman’ (in a particular context).¹⁶ This claim is far less metaphysically robust, in a certain sense, and so surely deflationists and realists in this literature don’t face the same challenge when it comes to bridging the gap between their normative and ontological commitments.

However, even if it’s true that this challenge is in some sense more pressing for the realist (given that their ontological commitments are somewhat “weightier”), this challenge is still present for the deflationist. For all this would mean is that, for deflationists, *weak* ontological commitments follow from normative commitments, and that’s not as alarming as it is for realists. But insofar as our question is one regarding how *any* claims about what is the case (weak or strong!) might follow from claims about what should be the case, this is unhelpful. A weak ontological claim is still an ontological claim.¹⁷ And so, bearing these considerations in mind, I will proceed as though our central line of inquiry here remains open.

¹⁴ See fn. 2: Díaz-León (2018). See also Amie Thomasson, *Ontology Made Easy*, (Oxford: Oxford University Press).

¹⁵ See fn. 12: Barnes (2017).

¹⁶ See fn. 2: Díaz-León (2018).

¹⁷ I will emphasize this at a few points throughout my discussion below, but it is crucial here that we are exploring a move from normative claims to *ontological* ones—not merely conceptual or linguistic ones. I say this because my sense is that some of the same philosophers who are sympathetic to metaphysical deflationism may also find it tempting to cash out these debates in terms of the notion of metalinguistic negotiation. These are negotiations which “concern the question of how a particular word should be used, which can include questions such as which concept should be paired with that word” (Plunkett & Sundell forthcoming, p. 1). Arguably, sometimes feminist metaphysicians are engaged in metalinguistic negotiation over gender terms and concepts. My project here does not deny that. But I’m also not interested in that. Here, what we want to know is how metaphysical claims about which entities exist might proceed from underlying normative commitments to justice. See Alexis Burgess, Herman Cappelen, and

4 “Ought Implies Is”

Everything so far has been by way of isolating and motivating this question regarding how the normative commitments underlying projects in feminist metaphysics might be somehow delivering the ontological commitments of these projects. I’m going to engage this question directly now, on the assumption that we’ve got space to ask it.

In particular, I’ll explore three ways that this relationship—between what should be the case and what is the case—might hold in this context: first, that normative commitments generate epistemic constraints on theories in feminist metaphysics (§4.1); second, that gender terms like ‘woman’ turn out to be normative, in much the same way as terms like ‘good’ and ‘just’ (§4.2); and third, that the normative commitments of these projects somehow directly deliver their ontological posits (§4.3). I will call these the epistemic interpretation, the metaethical interpretation, and the ontological interpretation of “ought implies is,” respectively.

I will hold this discussion in the context of inquiries into the nature of gender, as those inquiries are conducted in feminist metaphysics. And notably, a significant amount of that work backgrounds a normative commitment to the trans-inclusivity of gender, a normative commitment which then informs an ontological commitment to the existence of trans-inclusive gender kinds.¹⁸ In other words, where the question under discussion is something like:

QUD: What is it to be a woman?

These projects work with a particular normative commitment:

Normative commitment: the gender kind *woman* should be trans-inclusive.

David Plunkett, *Conceptual Engineering and Conceptual Ethics* (Oxford: Oxford University Press 2020); David Plunkett & Tim Sundell, “Metalinguistic Negotiation and Matters of Language: A Response to Cappelen,” *Inquiry: An Interdisciplinary Journal of Philosophy* (forthcoming): 1-25.

¹⁸ See for example fn. 2: Esa Díaz-León (2016). See also, fn. 3: Jenkins (2016). See also Robin Dembroff, “Real Talk on the Metaphysics of Gender,” *Philosophical Topics*, 46, 2 (2018): 21-50; Katharine Jenkins, *Ontology and Oppression: Race, Gender, and Social Reality*, (New York: Oxford University Press 2023).

And this normative commitment informs the eventual ontological posits of these projects:

Ontological commitment: the gender kind *woman* is trans-inclusive.¹⁹

In each of the sections below, I'll consider how different interpretations of "ought implies is" could operate in the background of philosophical investigations which can be generally characterized by the normative and ontological commitments above.²⁰

Of particular interest here is whether a given interpretation of the principle "*fits*" the context of inquiry. That is, does that interpretation of the principle fit the role feminist metaphysicians need it to perform here? And if a given interpretation "fits," we also need it to be *defensible*. That is, it must hold up under scrutiny. Here, I'll treat these like success conditions for an interpretation of "ought implies is." What we want is an interpretation which satisfies both of these conditions: an interpretation which fits the role set out for it and holds up under scrutiny. If a particular interpretation does not deliver on one or both of these counts, it fails.

A Brief Comment on the Nature of Normative Commitments

So far, I've been using the term 'normative commitment' in a very general manner. However, given the subject matter of this paper, it's more than fair to ask what these commitments are,

¹⁹ This investigation concerns the move from normative commitments to ontological ones, where the latter are commitments to the existence of certain metaphysical entities. Here, I'll discuss ontological commitments to gender *kinds*, but that I've chosen to talk in terms of kinds here inconsequential—the ontological commitment could be to properties, or systems, or relations, or groups . . . All that matters is that this is a commitment to the existence of a non-linguistic, worldly entity (i.e., an ontological commitment). Relatedly, in this section I'll try to remain as neutral as possible regarding the general metaphysical nature of gender kinds. If it helps, you can loosely think about these things as collections of individuals which share one or more properties in common. The specific characteristics of kinds matters less here than does the fact that these are non-linguistic, metaphysical entities.

²⁰ Of course, our working example necessarily oversimplifies the literature in a few ways. However, here I continue to work with the observation set out in the introduction of this paper. This is the observation that, despite their differences, these projects share a general methodology in common, one according to which their ontological posits proceed from their underlying normative commitments. My aim in this section is to explore how that very general move could work within the context of discussions about the nature of trans-inclusive gender kinds. It will take some further work to apply this discussion to each individual project which participates in that general methodology.

exactly. That is, are they something like moral facts, rules, or principals to which these theories are committed? And if so, in virtue of what do these commitments hold? We might also ask after their metaethical status: are we cognitivists or non-cognitivists about the content of moral claims here? Naturalists or non-naturalists about the nature of moral properties? Constructivists about morality, or objectivists? My sense is that the projects in feminist metaphysics discussed here don't really take a stance on these questions, and if that's right, I think this is actually quite unfortunate. For, given the authoritative role that normative commitments play in these projects, the nature of these commitments really matters.

In addition, given that the literature has largely been silent on the nature of these normative commitments, this also limits discussion in the present section somewhat, since it's difficult to analyze how normative commitments might inform ontological ones when we don't really know what the former commitments are, exactly. However, I do think we know enough to be able to get this inquiry off the ground. For example, we know that several projects in feminist metaphysics work with a trans-inclusive normative desideratum;²¹ we have Haslanger's list of normative, theoretical purposes which direct her ameliorative project;²² and we have arguments from philosophers like Dembroff and Jenkins' for the claim gender classification practices should be modeled on the classification practices indigenous to marginalized communities.²³ These examples give us something to work from as we explore this connection between normative and ontological commitments below. I maintain that I wish we had a bit more here, but my hope is that this will be enough for now.

4.1 Ought Implies Is: An Epistemic Interpretation

On the first way of understanding the relationship in question, the normative commitments of these projects somehow translate into *epistemic constraints* on theories in feminist metaphysics. And it is in their role as epistemic constraints on theories that normative commitments then influence the ontological posits of these projects. Let's call this the epistemic interpretation of "ought implies is:"

²¹ See fn. 18

²² See fn. 1: Haslanger (2012), 224-238.

²³ See fn. 18, especially Dembroff (2018) and Jenkins (2023).

EPISTEMIC The fact that P should be the case epistemically
constrains our theories to just those which posit P.

Considered within the context of our working example, facts about what gender kinds should be like serve as epistemic constraints on theories of gender:

EPISTEMIC The fact that gender kinds should be trans-inclusive
(*EXAMPLE*) epistemically constrains our theories of gender to just
those which posit trans-inclusive gender kinds.

Our task is now to determine whether *EPISTEMIC* can play a suitable and defensible role in projects which aim to answer our example question, regarding the nature of the gender kind *woman*. That is, how (if at all) could this version of “ought implies is” play a role in philosophical projects which seek to answer questions like “what is it to be a woman?”. We want to know whether *EPISTEMIC* fits this context of inquiry, and if it does, we’ll then want to determine whether this interpretation of “ought implies is” holds up under scrutiny.

Beginning with the first question, regarding fit: when we ask whether *EPISTEMIC* fits the context of inquiry, we’re asking whether this principle is positioned to take us from normative commitments regarding what should be the case, to ontological commitments regarding what is the case. It’s important to ask whether each interpretation of “ought implies is” fits the task set before it, otherwise we risk wasting our time exploring versions of the principle which are entirely inapplicable to projects in feminist metaphysics. So, applying *EPISTEMIC* within the context of our working example, a normative commitment to the trans-inclusivity of gender somehow serves to epistemically constrain our theories to just those which posit trans-inclusive gender kinds. This could happen in a few ways.

For example, perhaps empirical data concerning the oppression and marginalization of trans people gives us reason to construct theories of gender which can explain why this treatment is unjust.²⁴ Here, a normative commitment to the claim that gender kinds should be trans-inclusive

²⁴ See fn. 18, especially Jenkins (2016, 2023).

influences us to treat certain empirical data as evidence which plays a particular role in constraining our theories to just those with certain explanatory credentials. But this isn't quite strong enough. For, it's unclear that a theory of gender which accomplishes this explanatory work would need to posit trans-inclusive gender kinds. In fact, one might think that theories of gender which posit trans-*exclusive* gender kinds would fare at least as well here, in reflecting the harsh reality of the social world and thereby illuminating the ways in which it is unjust.²⁵ In other words, to theorize in this way isn't really to proceed in a way which backgrounds *EPISTEMIC*, because *EPISTEMIC* requires our theories of gender to posit trans-inclusive kinds (and the strategy sketched here leaves open that the resulting theories might not posit kinds like that).

We could avoid this outcome if our normative commitments worked like constraints on expressive and explanatory adequacy for theories of gender. In this role, they could serve to rule out theories which posit trans-exclusive gender kinds as inadequate. For example, if we begin here with a pre-theoretical commitment to the idea that the metaphysical reality of gender is trans-inclusive, this then generates constraints on expressive and explanatory adequacy for theories of gender, such that if a particular theory cannot capture the trans-inclusive reality of gender, that's a mark against the theory. But note that this approach also doesn't background *EPISTEMIC*. This is because we're *beginning* with an ontological commitment to the trans-inclusivity of gender kinds, and that commitment then gives rise to our normative commitments. In other words, this methodology exactly reverses the order of operations, for we're essentially just assuming that gender kinds are trans-inclusive and working from there. Of course, this is one way to proceed in theorizing about gender. However, this approach does not shed light on our question regarding how we might move from claims about what should be the case, to conclusions about what is the case.

Another way to go here could be to consider the potential effects of a particular theory of gender in practice. For example, it's plausible that the consequences of a trans-exclusive theory of gender in practice would be devastating for a great many people. We might then treat that data as evidence which weighs against theories of gender which posit trans-exclusive gender kinds.²⁶

²⁵ See fn. 1: Haslanger (2012). See also Ásta, *Categories We Live By* (New York: Oxford University Press 2018).

²⁶ See fn. 18, especially Díaz-León (2016) and Jenkins (2016).

The general idea here is that normative facts like [theory *T* would generate normatively bad consequences] and [theory *T* would generate normatively better consequences than would *T**] serve as evidence which can weigh in favor of/against theories of gender.²⁷ Elizabeth Barnes makes explicit elements of this methodology in Haslanger, pointing out that “[for] Haslanger, that a view has socially and politically unacceptable consequences is taken to be *evidence* of that view’s shortcomings.”²⁸

Within the context of our example, then: the idea is that a normative commitment to the claim that gender kinds should be trans-inclusive will take us to certain normative data concerning the effects of particular theories in practice; and, given that theories which posit trans-exclusive gender kinds will have normatively worse consequences than those which posit trans-inclusive gender kinds, this data serves as evidence which counts in favor of trans-inclusive theories of gender. As such, this way of implementing *EPISTEMIC* into our theorizing seems to fit the context of inquiry.

I’ll assume we have good reason to think *EPISTEMIC* fits the literature, given that at least some feminist metaphysicians have already been positioning their normative commitments in an epistemic role, in the ways discussed above. Now we can move on to our second task, which is one of determining whether this version of “ought implies is” holds up under scrutiny.

The discussion above makes salient at least two important questions: first, how exactly does a commitment to the claim that gender kinds should be trans-inclusive work to restrict the space of evidence which bears on our theories? And second, how exactly do normative data bear on the truth of our theories, more generally? To put a bit of pressure on that second question: even if we accept that some theories of gender have better consequences (in a normative sense) than others, and that given this, we ought to follow our consequentialist intuitions and go with whichever theory of gender produces the best consequences (perhaps in whichever sense is most relevant, given our aims), this still leaves us with the question of *how the normative consequences*

²⁷ See fn. 2: Díaz-León (2018; 2021).

²⁸ See fn. 12: Barnes (2017), 5-6.

of a theory might directly bear on its truth. That is, why should we think that one theory is more likely to be *true* than another, given that it would have better consequences?

These are significant questions, each of which poses a challenge to the viability of *EPISTEMIC*. For the overarching task here seems to be one of accounting for how the demands of justice might also be truth-tracking. That is, it seems we need to show that our normative commitments regarding what should be the case are also, somehow, a reliable guide to what is the case—a reliable guide to what’s true.²⁹

There are a few resources we might try to make use of here, perhaps most saliently from the literature in epistemology on moral encroachment. According to one of the more radical views in that literature, moral considerations can either (i) directly provide evidence for or against beliefs, or (ii) constrain epistemic norms so as to rule out otherwise rational beliefs.³⁰ Unfortunately, my sense is that the same challenges raised by the present discussion of *EPISTEMIC* are shared by this view in the moral encroachment literature. That is, the required connection between what’s right or good, and what’s true, is very difficult to establish. There may be other strategies that literature makes available, however, and so perhaps that gives us some reason to remain optimistic about the possibility of uncovering the relationship we need here, between truth and justice.³¹

But importantly, we are in need something stronger than the claim that the demands of justice are possibly truth-tracking. I’ve proposed *EPISTEMIC* as a potential interpretation of “ought implies is,” a principle which represents one of *the central methodological moves* in feminist metaphysics. A defense of *EPISTEMIC* requires an account of this relationship between justice and truth, one which explains how facts about what gender should be like can take us to facts about the way gender is. Without this, *EPISTEMIC* fails.

²⁹ Framing things this way brings out that, even if the normative considerations at play in feminist metaphysics are not consequentialist in nature, there is still a challenge here. That is, even if there is a way to reformulate the normative data which are central to these projects, we may still face a more general version of this same worry. And that is: how might normative constraints *of any sort* translate into epistemic constraints which bear on the truth of a theory?

³⁰ See Renée Bolinger, “Varieties of Moral Encroachment,” *Philosophical Perspectives* 34, 1, (2020): 5-26; Rima Basu, “The wrongs of racist beliefs,” *Philosophical Studies*, 176, 9 (2019): 2497-2515.

³¹ Thanks to [redacted] for incredibly helpful discussion here.

4.2 Ought Implies Is: A Metaethical Interpretation

It might seem like this gap between normative and ontological commitments is going to be exceptionally difficult to bridge. If that’s right, perhaps it would be better if there were no gap at all. In this section, I’ll explore one way of arguing that, despite appearances, this there is no gap here between the normative and the non-normative; instead, there is normativity “all the way up,” in a certain sense.

More specifically, on this second way of understanding the relationship in question, the ontological posits of theories in feminist metaphysics follow from underlying normative commitments insofar as the central terms used to articulate these ontological claims (e.g., ‘woman’, ‘man’) turn out to be *normative terms* which express normative properties, in much the same way as terms like ‘good’ and ‘just’. Let’s call this the metaethical interpretation of “ought implies is”:

METAETHICAL From the fact that P should be the case, it follows that
P is the case (where ‘P’ is a normative term, and so [P
is the case] is a normative claim).

Considered within the context of our working example, this version of “ought implies is” has it that normative facts about what gender kinds are like follow from normative facts about what should be true of those kinds:

METAETHICAL From the fact that the gender kind *woman* should be
(*EXAMPLE*) trans-inclusive, it follows that *woman* is trans-inclusive
(where ‘woman’ is a normative term, and so [*woman* is
trans-inclusive] is a normative claim).

Just as in the last section, our task is now to determine whether *METAETHICAL* can play a suitable and defensible role in projects which aim to answer our example question, regarding the nature of gender kinds like *woman*. We want to know whether *METAETHICAL* fits this context of inquiry, and if it does, we’ll then want to determine whether this interpretation of “ought implies is” holds up under scrutiny.

This discussion takes us away from epistemology and into the territory of metaethics. Here, we're faced with number of substantive choice points right off the bat, as there are a number of ways to interpret the claim that gender terms like 'woman' are normative terms. Among other things, this means there are many specific ways that the general strategy in this section could be carried out, and I think that's exciting. However, given constraints on space, I will explore a very narrow range of those options here. Specifically, I'll proceed with the background view that moral claims express propositions, and normative terms (elements of those propositions) express normative properties. So when I say, "Giving to charity is good!" that claim expresses the proposition that it is good to give to charity, where the word 'good' is a normative predicate which expresses a normative property, being good.

Now the suggestion to be explored in this section can be put as follows: claims like "Tituba was a woman" are moral claims which express propositions (e.g., [Tituba was a woman]); the word 'woman' expresses a property, being a woman, where that is a normative property, similar to the property of being good.

Now, why think gender properties like this might be normative? This question is best answered in terms of what might make any property a normative one. Here I will discuss a general version of one dominant view in the literature on this subject: namely, the view that normativity is to be analyzed in terms of *reasons*. On this view, paradigmatically normative properties like being good and being just are analyzed in terms of reasons, such that e.g., the fact that an action is good or right counts as a reason in favor of performing it; the fact that a state of affairs is just counts as a reason to bring it about; or the fact that something ought to be the case counts as a reason to make it so.³² Versions of this general view have been widely discussed and defended

³² Very generally, we can say R is a reason to Φ just in case R counts in favor of Φ -ing. Note that some people think facts like e.g., [action A is morally right] are themselves reasons to perform A, while others think that A's being morally right is *not* a reason to perform A (but rather whatever makes A right is the source of our reasons here). For the sake of simplicity, my examples above reflect the first way of thinking about things, but for the purposes of this section, it really doesn't matter whether we should "pass the buck" on the reasons-giving nature of moral rightness, or goodness, etc.

in moral philosophy, and so I'll proceed here with the assumption that this is a plausible analysis of what it takes for a property to be normative.³³

Are gender properties normative in this sense? Of course, the answer here will depend largely on what gender properties are like. The literature on the metaphysics of gender offers us a range of views, and here I'll work with two dominant ways of theorizing about the nature of gender properties: social position views and self-identification views.³⁴ Now we can refine our question here a bit: we want to know whether gender properties, like the property of being a woman, are normative (in the reasons-giving sense) on either social position views of gender, or self-identification views of gender.

Very generally, according to social position views, whether someone is e.g., a woman is determined largely by how other people and institutions treat them in their context. Most (if not all) social position views are designed to bring out the relational nature of gender, and many times the relevant relations are power relations. Sally Haslanger has developed one of the best-known social position accounts of gender. Very generally, on her view, someone is a woman just in case they are systematically subordinated because they are taken to be biologically female. In this way, Haslanger's metaphysics of gender essentially incorporates elements of oppression and subordination, and importantly, these are normative notions. For Haslanger, that there are women is not a good thing, for this fact is made true by deeply unjust states of affairs. And she argues that, as feminists, our aim should be to eliminate that injustice, where success in this regard will have the consequence that there are no longer any women or men. A just world, on Haslanger's view, is a genderless one.

³³ See for example T.M. Scanlon, *What We Owe to Each Other*, (Cambridge: Harvard University Press 1998); Mark Schroeder, "Realism and Reduction: the Quest for Robustness," *Philosophers' Imprint*, 5, 1 (2005): 1-18; Joseph Raz, *Explaining Normativity: On Rationality and the Justification of Reason*, reprinted in *Engaging Reason*, (Oxford: Oxford University Press 1999), 67-89.

³⁴ For examples of social position views of gender see fn. 25 above. For examples of self-identification views of gender see Talia Mae Bettcher, "Trans Identities and First-Person Authority" In Laurie Shrage (ed.), *You've Changed: Sex Reassignment and Personal Identity* (Oxford: Oxford University Press 2009); Jennifer McKittrick, "A dispositional account of gender," *Philosophical Studies*, 172, 10 (2015): 2575-2589; Susan Stryker & Stephen Whittle (ed.), *The Transgender Studies Reader* (NY: Routledge 2006).

So defined, is the property of being a woman a normative property, in the sense that it is reasons-giving? Plausibly so. To see this, consider the fact that there are women. On social position views like Haslanger's, this is equivalent to the fact that there are people who are unjustly subordinated because they are taken to have female biology. I think it's clear that Haslanger herself takes this fact to give us reasons for action, given the shape of her overall ameliorative project. In a certain sense, the whole *point* here is to characterize gender in a way that motivates political change—in a way that gives us reasons to act.³⁵

It's plausible that gender properties can also be normative on many self-identification views. Very generally, according to self-identification views of gender, whether someone is e.g., a woman is given by facts about their psychology, and how they understand themselves in relation to their immediate environment, and the broader social world. Could gender properties be normative on self-identification views of gender, in this reasons-giving sense? Again, plausibly so. Several self-identification theorists have argued that facts about the way someone self-identifies generate reasons for us to treat them in certain ways which respect that self-identification.³⁶ For example, if someone self-identifies as a woman, that gives us reasons to treat them in accordance with the way they identify. Of course, this can and does vary between individuals, but for many women, this can mean things like e.g., applying more feminine-coded gender norms to them, using 'she/her' pronouns to talk about them, not referring to them as men, etc. Notably, these reasons may not *ultimately* be rooted in their gender identification (they may instead be ultimately grounded in something like e.g., respect for persons), but on these views, someone's gender self-identification is integral to the explanation for the reasons we have to regard them and treat them in certain ways. If that's right, then self-identification gender properties can be reasons-giving.

³⁵ Importantly, not all social position accounts have it that gender is essentially hierarchical in this way. For example, Ásta has developed another prominent social position analysis of gender as a conferred property. On her view, gender properties can be, and often are, conferred under unjust and oppressive conditions, but this is not necessarily the case. So, in applying our discussion about normative properties to views like Ásta's it is important to account for the specifics of different modes of gender conferral; it may be the case that not all gender properties are normative on Ásta's view, but plausibly some of them will be. See fn. 25: Ásta (2018).

³⁶ See fn. 34, especially Bettcher (2009) and Stryker & Whittle (2006). See also Robin Dembroff & Daniel Wodak, "He/She/They/Ze," *Ergo*, 5, 14 (2018): 371-406. Quill Kukla & Mark Lance, "Telling Gender: the Ethics and Pragmatics of Gender Ascriptions," *Ergo* (forthcoming).

Now according to *METAETHICAL*, facts like [the gender kind *woman* is trans-inclusive] follow from facts like [the gender kind *woman* should be trans-inclusive], where both sorts of facts are normative. And this version of the principle seems to fit the literature, insofar as it takes us from claims about what should be the case to claims about what is the case. And, on several existing views of gender in feminist metaphysics, it's plausible that gender properties can be normative. If that's right, then we can move on to our second task, which is one of determining whether this version of "ought implies is" holds up under scrutiny.

The general strategy behind *METAETHICAL* involves illuminating that the ontological claims here are also normative. And so, contrary to appearances, this principle represents a move from a normative claim to another normative claim. As such, perhaps *METAETHICAL* is just like other familiar normative principles which take us from observations about how we should act, or what should be the case, to conclusions about the way the world is—the sorts of rights and duties people have, which things are virtuous and vicious, the normative properties characteristic of personhood, etc.

But hold on a moment. *Is* that how those familiar normative principles work? Consider, for example, the Kantian notion of respect for persons. According to Kant, I shouldn't treat my neighbor (or any other person) as a mere means to my ends, and this is because there is something normatively special about personhood which makes that sort of treatment wrong. In other words, I ought to respect my neighbor *because* he is a person. This fact about my duties (that I ought to respect my neighbor) and this fact about my neighbor (that he is a person) are both normative. But importantly, in this case, the claim about what I should do follows from the descriptive claim about my neighbor's personhood, not the other way around. This is exactly the reverse of *METAETHICAL*.

According to *METAETHICAL* (applied in the context of our working example), the gender kinds are trans-inclusive because they should be. That is, the descriptive claim follows from the should claim. If Kant's principle worked in that direction, it would have the consequence that my neighbor is a person *because* I should respect him. But that's wrong: my duties and obligations to other people are grounded in the sorts of beings they are, not the other way around. So at

least on the face of it, it won't work to defend *METAETHICAL* by arguing that it is just like other familiar principles, according to which normative claims generate further normative claims.

More generally, this discussion brings out that the move from normative commitments to ontological ones isn't puzzling because it is a move from a normative claim to a non-normative one. It's puzzling because it is a move from a claim about way the world should be, to a claim about the way the world is, where the ontological claim seems to be generated by and strongly dependent upon the normative claim. *METAETHICAL* doesn't shed any light on that relationship however, or how that connection might hold. And so, for at least this reason, *METAETHICAL* doesn't leave us with a defensible version of "ought implies is."

4.3 Ought Implies Is: An Ontological Interpretation

There is one final option to be considered here, one which amounts to a literal interpretation of "ought implies is" in this context. I'll call this the ontological interpretation of the principle:

ONTOLOGICAL From the fact that P should be the case, it follows that P
is the case.

Considered within the context of our working example, this version of "ought implies is" has it that facts about what gender kinds are like follow straightforwardly from facts about what should be true of those kinds:

ONTOLOGICAL From the fact that the gender kind *woman* should be
(*EXAMPLE*) trans-inclusive, it follows that *woman* is trans-inclusive.

I've left this discussion for last because I have the least to say about this version of the principle by way of either an analysis or a defense. For, in a certain sense, *ONTOLOGICAL* asserts that the connection we've been investigating *simply obtains*. And this is one way to go here: we might just assume that, at least sometimes, ontological claims about what is the case follow from normative claims about what should be the case. And so, perhaps this is best understood as just a basic

assumption at the root of our theorizing in feminist metaphysics, and there is little more to be said about this principle, apart from remarking on its presence and accounting for its role.

If you find this unsatisfying (as I do), there are two more options available here, both of which are suggested by Anderson.³⁷ For, although she thinks that debates over “ought implies is” ultimately turn out to be a red herring, she does offer two points in defense of an ontological reading of this principle.

First, she suggests that *ONTOLOGICAL* might hold in virtue of a more fundamental teleological law of the form “If P ought to be the case, then P is the case”.³⁸ As such, if proponents of value-neutral approaches to theory choice aim to argue against “ought implies is,” they are tasked with the project of showing that the universe is not governed by teleological laws like this.

My project here is not one of defending value-neutral approaches to theory choice, and so I won’t attempt to undertake the task Anderson suggests. However, my project here *is* one of seeking out a plausible interpretation and defense of “ought implies is” and with that project in mind, this suggestion from Anderson is not particularly helpful. I grant that it is logically *possible* that teleological laws like this govern the universe, but given the importance of our inquiry here, it would be nice if there were something more to say than simply that it is not impossible that there exists a law which grounds this principle. For, contra Anderson, the assumption that there are no laws like this doesn’t strike me as controversial in the slightest (largely because I don’t know of any contemporary metaphysical worldviews which posit laws of this form).

Anderson’s second suggestion is essentially that “ought implies is” works similarly to “ought implies can.” And so, insofar as the latter is a plausible principle, so too is the former.³⁹ However, the success of this argument seems to turn on a rather unorthodox interpretation of the principle of “ought implies can.” To see this, consider: on the ontological interpretation of “ought implies is,” facts like [P should be the case] somehow *make true* or otherwise generate facts like [P is the case]. An analogous reading of “ought implies can” would have it that facts like [S ought to

³⁷ See fn. 9: Anderson (1995a).

³⁸ See fn. 9: (Anderson 1995a), 34.

³⁹ See fn. 9: Anderson (1995a), 34.

perform A] make true, or otherwise generate facts like [S can perform A]. But this is not how the latter principle is typically understood. Instead, those who endorse the principle of “ought implies can” typically take something like the following to be true: that S ought to perform A implies that S can perform A, such that if S cannot perform A it follows that it is not the case that S ought to perform A.⁴⁰ But this is very different from the claim that S can do A *because* S ought to do A! And yet, this is what is required by “ought implies is.”

But perhaps Anderson is suggesting that we read “ought implies is” in terms of the orthodox understanding of “ought implies can.” That is: the fact that P should be the case implies that P is the case, such that if P is not the case, then it is not true that P should be the case. But note that this reading is hardly something that feminist metaphysicians will find attractive. For on the assumption that the normative commitments which underly these projects do not *necessitate* their ontological outputs, it is possible that these normative claims hold while the corresponding ontological claims do not. And this is desirable. For consider: if we were to endorse the present reading of “ought implies is” in the context of our working example, it would mean that if gender kinds do not turn out to be trans-inclusive, it follows that it’s not the case that they should be! I am doubtful that anyone in the literature would endorse this. And so, if this were the most plausible reading of “ought implies is,” that would be an unhappy outcome for many feminist metaphysicians.

In conclusion: if we interpret ONTOLOGICAL in terms of the orthodox understanding of “ought implies can,” it does not fit the context of inquiry. And if we interpret ONTOLOGICAL in terms of an unorthodox understanding of “ought implies can,” the principle is opaque. In other words, ONTOLOGICAL either doesn’t fit the demands of the literature, or it does fit and it’s indefensible. Either way, this interpretation of the principle fails.

5 Conclusion: Feminist Philosophy & the Aims of Metaphysical Inquiry

Let’s sum up the discussion so far. Projects in feminist metaphysics seem to background a principle of “ought implies is,” insofar as the ontological posits of these projects somehow proceed from their underlying normative commitments. Our task has been one of providing an

⁴⁰ Brian Talbot, “The Best Argument for ‘Ought Implies Can’ Is a Better Argument Against ‘Ought Implies Can,’” *Ergo*, 3, 14 (2016): 337-402.

analysis and defense of that connection between what should be the case and what is the case, as this relationship has been posited by projects in feminist metaphysics. To this end, we've explored three interpretations of the principle:

<i>EPISTEMIC</i>	The fact that P should be the case epistemically constrains our theories to just those which posit P.
<i>METAETHICAL</i>	From the fact that P should be the case, it follows that P is the case (where 'P' is a normative term, and so [P is the case] is a normative claim).
<i>ONTOLOGICAL</i>	From the fact that P should be the case, it follows that P is the case.

Each of these interpretations faced certain challenges. Of course, there may yet be additional ways to interpret and defend this principle which haven't been considered here. For example, it could be worthwhile to return to the epistemic interpretation and invoke resources from the literature on moral encroachment to analyze the manner in which normative constraints might translate into epistemic constraints on theories (§4.1). Alternatively, we might try to find an additional interpretation of the principle (perhaps entirely different from those I've offered here) and defend that. And of course, we might also just abandon the project of providing an analysis and defense of "ought implies is," and simply treat the principle like a basic axiom at work in feminist theorizing (§4.3).

In closing this essay, I'd like to return to one final strategy, something suggested by discussion of Anderson at the outset of this paper. This was the suggestion that feminist metaphysicians might be pragmatists about the nature of (social) truth, and so avoid the challenge posed by "ought implies is" by formulating their ontological posits in terms of commitments to the pragmatic utility of certain claims (§3).

The first thing to note here is that this would be a quite radical conclusion. That is, if endorsing this form of pragmatism is the most viable way for feminist metaphysicians to maintain the

connection between the normative and ontological commitments at work in our projects, that is significant. It would also be surprising, given that many philosophers working in this literature do not explicitly endorse pragmatism. And it would be a significant point of departure from contemporary projects in more traditional “mainstream” metaphysics, projects which typically background a correspondence theory of truth, and which largely take truth (thus understood) to be the primary aim of theoretical inquiry.

I expect that some feminist metaphysicians will be content with this pragmatist outcome, however. For, why should we think that projects in feminist metaphysics must share the same theoretical aims as those of projects in the more traditional analytic literature? Perhaps this is just one among many things which sets our discipline apart, and so it would be a mistake to evaluate projects in feminist metaphysics terms of the methodology of the more traditional literatures.

I’d like to push back against this conclusion for a moment, however. For, even if it turns out that *some* projects in feminist metaphysics work well against a pragmatist backdrop, I think it would be unfortunate if it turned out that the only way to make sense of the dominant methodology in this discipline were in pragmatist terms. Very generally, this is because many feminist metaphysicians take themselves to be engaged in straightforwardly descriptive projects, where the aims of those projects are not just to produce e.g., helpful or especially useful ways of representing the social world, but theories which accurately reflect the nature of social reality.

As just one example of this: I take it that recent critiques of trans-exclusive theories of gender target not only the *utility* of theories which identify sex with gender, but also their objective *accuracy*. A view according to which women are adult human females *gets the world wrong*. Now, it’s also true that these trans-exclusive theories of gender get the world wrong in a way that matters to many people (and so pragmatic utility has a role play here, as well). But importantly, the critique of these views is not that it is unhelpful to describe e.g., trans women as men; the

objection is that to do this is to misrepresent the way the social world is, and the way people are.⁴¹

But importantly, to theorize from this descriptive angle is not to cut the normativity out of feminist theorizing all together. In fact, several feminist metaphysicians have argued that this sort of straight-forwardly descriptive project is integral to working towards justice.⁴² The general idea is that, in order to bring about a just world, we need to be able to identify, account for, and understand the myriad ways in which the actual world is unjust. This is the role of descriptive projects in feminist metaphysics: to accurately reflect the nature of the world as it is right now, so that we can change it for the better. But of course, in order to do this, we need a descriptive theory the content of which *corresponds* to the way the world is. And so, we cannot be pragmatists here.⁴³

Ultimately, this may leave those of us working in feminist metaphysics with a choice. On the one hand, we can be pragmatists about (social) truth. If we go this route, it's easy to see how our ontological claims about what is the case might follow from our normative claims concerning what should be the case, and so we can maintain a version of "ought implies is." And insofar as this principle lies at the heart of many projects in feminist metaphysics, this is a desirable outcome. However, it would also mean giving up on the sort of straightforwardly descriptive projects described above. And if it's right that these descriptive projects play an integral role in the overarching feminist project of working toward a just world, this would be a significant loss. And so, on the other hand, we might think these descriptive projects are too

⁴¹ See for example Robin Dembroff, "Escaping the Natural Attitude About Gender," *Philosophical Studies*, 178, 3 (2021): 983-1003; Rebecca Mason, "Women are not adult human females," *Australasian Journal of Philosophy* (forthcoming).

⁴² See fn. 25: Ásta 2018. See also fn. 18: Jenkins (2023). See also Elizabeth Barnes, "Feminist Metametaphysics," (forthcoming); Sally Haslanger, "Social construction: the debunking" project," in *Socializing Metaphysics: the Nature of Social Reality* (Rowman & Littlefield 2003): 301-325.

⁴³ Or at least, we cannot endorse the version of pragmatism I've discussed here, which seems to be required when it comes to closing the gap between our normative and ontological commitments (§3). Importantly though, there may be a number of ways to combine a pragmatist methodology with work in feminist metaphysics. Exploring those options could take us to another viable way of answering the challenge posed by "ought implies is," and I think this could be a promising project to pursue in future work on this subject.

valuable to abandon. However, insofar as these projects are incompatible with the sort of pragmatism described here, we cannot then avail ourselves of that particular solution to the challenge posed by “ought implies is.” We may yet have other options here, as I’ve suggested above.⁴⁴ But it remains an open question whether any of those strategies will succeed.

Acknowledgements

My deepest thanks go out to Matthew Andler, Elizabeth Barnes, Rebecca Chan, Jade Fletcher, Carolina Flores, Theodore Locke, Trenton Merricks, Asya Passinsky, Kevin Richardson, Matthew Duncan, the members of the Georgetown Philosophy Department, and the participants at the 2nd SoCal Metaphysics Network Conference and the 4th Annual Social Metaphysics Workshop for their many helpful contributions to this project.

⁴⁴ It might seem like this dilemma that I’ve raised here could be avoided by adopting something like a hybrid approach to theorizing about gender. That is, we could conduct a straight-forwardly descriptive project, one which does not background “ought implies is,” in order to get the lay of the land. And then we could engage in a normative project, one guided by our commitments to how the world should be. Jenkins (2023) articulates a version of this approach, and I personally find this strategy attractive. However, note that this strategy amounts to abandoning “ought implies is,” for at no point are our normative commitments taking us to claims about what is the case. Instead, we are considering questions about what is and what should be under the umbrellas of separate, but related, projects.