

9

Descriptive vs. Ameliorative Projects

The Role of Normative Considerations

E. Díaz León

1. Descriptive and Ameliorative Projects

Sally Haslanger (2000, 2006) has distinguished between *descriptive* projects and *ameliorative* projects in philosophy. The main idea is this: philosophers engaged in a descriptive project aim to reveal the *operative* concept, that is, the objective type that our usage of a certain term tracks (if any), whereas philosophers engaged in an ameliorative project aim to reveal the *target* concept, that is, the concept that we should be using, given our purposes and goals in that inquiry. The questions pertaining to ameliorative projects are the following: what is the point of having this concept? Which concept would serve these purposes best?

Haslanger (2000) offered a social constructivist analysis of gender and race, and argued that those analyses are not intended to capture our ordinary concepts of gender and race, but her aim was rather to figure out the target concepts, that is, the concepts of gender and race that would be most useful in order to achieve social justice. These analyses go as follows:¹

A group *G* is a *gender* (in context *C*) iff_{df} its members are similarly positioned as along some social dimension (economic, political, legal, social, etc.) (in *C*) and the members are “marked” as appropriately in this position by observed or imagined bodily features presumed to be evidence of reproductive capacities or function.

A group *G* is *racialized* (in context *C*) iff_{df} its members are similarly positioned as along some social dimension (economic, political, legal, social, etc.) (in *C*), and the members are “marked” as appropriately in this position by observed or imagined bodily features presumed to be evidence of ancestral links to a certain geographical region. (Haslanger 2003: 8)

However, in more recent work (2005, 2006), she has argued that her social constructivist accounts of gender and race could also be seen as trying to capture the

¹ These formulations are a simplified version of her original, more complex characterization in her (2000) paper.

operative concept that we actually associate with our terms ‘gender’ and ‘race’. As she argues, the operative concept could differ from the *manifest* concept (i.e., the concept that people take themselves to be using). Crucially, the characterizations of gender and race that would come to mind if ordinary speakers are asked ‘What is race?’ or ‘What is gender?’ might differ from the objective types that our uses of the terms ‘gender’ and ‘race’ actually track. In other words: even if the manifest concepts that many people associate with the terms ‘gender’ and ‘race’ are not explicitly social constructivist, these terms might refer to socially constructed properties after all.

This important distinction gives rise to the following question: should philosophers of gender and race engage in the descriptive project, or in the ameliorative project? It could be argued that these are two independent projects and that they are both useful. But if they are both useful, we can still ask: useful for what purposes, and under what conditions? In this chapter I would like to discuss this important question. In particular, I want to examine the nature and the prospects of both the descriptive and the ameliorative project regarding philosophy of gender and race. My main questions will be the following: what does the descriptive project consist in, and what steps are required in order to successfully complete it? What are the main criteria in order to evaluate different answers to a descriptive project? Furthermore, what does the ameliorative project consist in? What are the main criteria in order to evaluate different answers to an ameliorative project? Furthermore, are there any important connections between these two projects? If we have found a satisfactory answer to the descriptive project, does that pose constraints to our possible answers to the ameliorative project? And likewise, if we have found a satisfactory answer to the ameliorative project, does that pose constraints to our possible answers to the descriptive project?

Mari Mikkola (2015) argues that whereas philosophers of science have long recognized the role played by contextual factors in our inquiries (i.e., “the political and moral values embedded in the social context of an inquiry” (p. 782)), contemporary metaphysicians often assume that contextual factors do not belong in metaphysics. One of my main aims in this chapter is to argue that moral and political considerations among other contextual factors are relevant in metaphysics, and more in particular, I will show that they can be relevant at different stages of a metaphysical inquiry. I will start by discussing the nature of descriptive projects seeking to reveal the operative concept associated with a term, and I will argue that moral and political considerations are relevant at different stages of this project. In addition, I will distinguish two kinds of ameliorative projects (namely, cases where a term determinately refers to an entity but we ought to change the referent, and cases where a term’s referent is indeterminate but it should become determinate, given normative considerations). I will argue that many projects that are usually taken to be purely descriptive are actually ameliorative or contain important ameliorative elements.

2. The Descriptive Project Revisited

Let’s start with the descriptive project. Haslanger (2005, 2006) argues that the operative concept associated with gender or race might be a social constructivist concept after all. That is, the objective type that our concept of gender (or race) tracks

might be a socially constructed property, like the ones she characterized above. What are the conditions for this proposal to be correct? That is, what would determine that our operative concepts of gender and race actually pick out socially constructed properties, and in particular the socially constructed properties that Haslanger (2000) characterizes?

Haslanger (2005, 2006) appeals to semantic externalism in order to defend this possibility. That is, she argues that our terms do not get their referents fixed in terms of the descriptions, conceptions, and beliefs that we associate with them, but rather in terms of externally individuated factors. Her strategy is to extend the familiar view of semantic externalism about natural kind terms to social kind terms as well. In particular, she characterizes externalism as follows:

Objective type externalism: Terms/concepts pick out an objective type, whether or not we can state conditions for membership in the type, by virtue of the fact that their meaning is determined by ostension of paradigms (or other means of reference-fixing) together with an implicit extension to things of the same type as the paradigms. (2006: 110)

If we apply objective type externalism to gender and race (as objective types), it then follows that our concepts of gender and race get their referents fixed by means of *ostension*, that is, our terms ‘gender’ and ‘race’ refer to whatever objective types unify paradigms of gender (and likewise for race). The idea here, I take it, is that competent speakers (or at least the original speakers who introduced the term) would be able to point to paradigm cases of genders or races (e.g., gender groups, or racial groups). The referent will then be determined as follows: ‘let “gender” refer to the most objective type that all those instances have in common’, and likewise for the term ‘race’. A first worry arises: how can we identify the most objective type that some entities share? As I understand Haslanger’s view, the property that unifies all paradigmatic instances of a putative social kind will be that property that can better explain the paradigmatic features of the kind. That is, the most objective kind corresponds to the kind that satisfies some sort of explanatory role.

However, it could still be argued that this account leaves room for indeterminacy, since it is not clear that there is always a unique objective type that unifies paradigmatic instances of genders or races, since paradigms of genders have many properties in common, and paradigms of race have many properties in common, so it will be difficult to figure out which unifying property is the referent.² The crucial question, then, is the following: how can we identify the most objective property, out of the many properties that are shared by those paradigmatic instances? Appealing to their explanatory role might not be enough to identify a unique property. Or the situation might be even worse for the advocates of social constructionism: it could be argued

² As I am formulating the issue, the relevant terms are ‘gender’ and ‘race’, where paradigmatic instances that fall under the former would be properties or kinds such as *men* and *women*, and paradigmatic instances that fall under the latter would be racial groups such as *Black*, *White*, *Asian*, and so on. The issue could also be formulated concerning terms for specific genders or racial groups, such as ‘woman’ or ‘Black’. In this case, paradigmatic instances falling under the first term would be individuals such as Hillary Clinton, whereas paradigmatic instances of the second would be individuals such as Barack Obama. The question then would be: which objective type, if any, unifies all paradigmatic instances falling under the term?

that the most objective, explanatory property that unifies our paradigmatic cases of gender (or race) is actually a biological property. Haslanger discusses this indeterminacy worry, and she says: "Sets of paradigms will typically fall within more than one type. To handle this, one may further specify the kind of type (type of liquid, type of artwork), or may (in the default?) count the common type with the highest degree of objectivity" (2006: 110). The suggestion here, as I understand it, is that our concepts of gender and race might involve some *sortal* information (i.e., information about the kind of thing something is) constraining the possible candidate referents, such as 'social kind that those instances have in common', or something along those lines. If this is indeed the sortal information associated with our concepts of gender and race, then the relevant objective type in the vicinity can be a social kind. However, this move seems to beg the question against the advocate of biological realism, who would claim that concepts of gender and/or race do not involve sortal information about their being social kinds. Perhaps the social constructionist could argue that the sortal information associated with terms such as 'gender' or 'race' is not explicitly social, but it does rule out biological properties as possible referents. My point here is that if the social constructivist wants to argue that gender and race concepts involve sortal information that rules out biological properties as referents, they would need to provide independent motivation for this.

Haslanger also suggests that in some cases there might be no sortal information associated with the concept, and then the referent will just be by default the most objective type shared by the paradigms, whatever that is. But then we face our original worry: there might be more than one objective type in the vicinity (all being equally explanatory), or even worse for the social constructivist, the most objective type might turn out to be a biological property after all. In my view, one way of avoiding this worry on behalf of social constructionism would be to allow that concepts of gender and race can be associated with some information (sortal or otherwise) that rules out biological kinds as possible referents, so that social kinds can be the referents after all.

In any case, in order to settle this dispute about what kind of information is associated with our concepts, one would need to engage in something very similar to what Haslanger (2006) calls the *conceptual* or *analytical* project. She characterizes this as the project of revealing our *manifest* concept, that is, the concept that we take ourselves to be using, or would come easily to mind if we are asked. However, I believe we should characterize the conceptual project in terms of the search for the *application conditions* of our concepts (which is a fallible inquiry, and which may require a lot of reflection about actual and possible cases, and even empirical research).³

³ Haslanger (2005, 2006) seems to conflate these two different characterizations of the *manifest* concept. On one reading, the manifest concept corresponds to the application conditions of the concept. On this interpretation, the manifest concept corresponds to the conditions that something should satisfy in order to fall under the term. For example, the manifest concept associated with 'water' would correspond to something like 'the actual watery stuff', whereas the operative concept would correspond to H₂O, since this is the property that satisfies the manifest concept in the actual world. On another reading of 'manifest' concept, this corresponds to the concept we take ourselves to be using with a certain term, which may or may not correspond to the way the term is actually used in that linguistic community or practice.

But if we agree that the conceptual project is necessary in order to avoid indeterminacy, then we would have to deny one of Haslanger (2006)'s central claims, namely, that the descriptive project does not require the conceptual project. In my view, this only requires a friendly modification of Haslanger's approach, as follows. It could be argued that when it comes to the descriptive project, we need to engage in a *two-step* process: first, we need to find out the core information that competent speakers associate with the concept (where this might be accessible relatively *a priori*, although it might involve a lot of careful considerations of actual and possible cases, and even experimental semantics); and, second, we need to find out what objective kinds, if any, satisfy those descriptions in the actual world (where this can be found out only empirically).⁴ However, Haslanger's main claim still applies: when it comes to figuring out the operative concepts of gender and race, that is, what objective properties our concepts of gender and race actually track, there is at least a central part of this inquiry that will be *a posteriori* or empirical, and social constructionism is still a live option.

In order to illustrate this two-step conception of the descriptive project, it will be useful to consider an example. Joshua Glasgow (2009) has argued that ordinary speakers associate the concept of race with the belief that racial groups are characterized in terms of certain visible traits that most members of a racial group have in common. He argues that ordinary speakers wouldn't be willing to give that belief up without replacing the ordinary concept of race with another concept. Glasgow (2009) then uses this claim about the application conditions of the concept of race in order to argue that the candidate meanings for 'race' that have been proposed by social constructivists and biological realists about race cannot actually be the referents of our ordinary concepts of race. For instance, he argues that reproductively isolated biological populations cannot be the referent of our ordinary concept of race, given that those biological populations do not satisfy some of the central features that are part of our ordinary concept of race (namely, that racial groups correspond to specific visible traits, whereas reproductively isolated populations do not need to have visible traits in common). Likewise, he argues, social structures of the sort advocated by Haslanger do not satisfy those central features of our ordinary concept either (because members of socially constructed groups do not have to have visible traits in common), and therefore social constructions cannot be the referent of 'race' either. This line of argument assumes that there is some information associated with an ordinary concept such that a candidate referent must satisfy it in order to be the referent. On the contrary, Haslanger (2005, 2006) explicitly denies this assumption (and suggests that the referent of 'race' is the most objective property by default). However, as I have suggested, I think that advocates of social constructionism about the operative concept of gender or race should not reject the first step of the descriptive project. In particular, this two-step approach can help the social constructivist to respond to Glasgow's objections to social constructivism, as follows.

For instance, Haslanger (2006) considers the example of a school that uses the term 'parent' in school memos but actually means 'primary caregiver' by it. Here I am focusing on the former characterization. See Díaz-León (2012) for further discussion of this distinction.

⁴ Here I draw on Jackson (1998); Chalmers and Jackson (2001); and Thomasson (2007, 2008).

It could be argued that given the information associated with 'race' and the paradigm cases of racial groups that ordinary speakers would point to, there are no biological properties that unify all those cases, because, for instance, there are no biologically significant properties unifying all individuals that we would call 'white', and that the only properties that unify these paradigm cases are social properties. But in order to formulate this argument, we need to rely on the premise that there is some information associated with the term 'race' that rules out biological properties as being the referent. In addition, as I explained above, the conceptual project could also help with the worry that if we do not appeal to any associated information at all, then the referents of 'gender' and 'race' might be indeterminate, for it could be argued that the only feasible way to avoid this indeterminacy would be to appeal to some central information associated with our concepts that puts some constraints on what shared properties can be the referents. On the other hand, if we renounced to the first step of the two-step descriptive project, as Haslanger does, then the prospects of finding a good response to Glasgow's mismatch objections to social constructionism about race look very dim.⁵

Once this is clarified, we can see that the social constructivist could in principle argue that social structures might turn out to be the most objective property that unifies paradigm cases (assuming that they satisfy the information associated with the concept, if any). But how could we establish that a socially constructed property is actually the most objective, explanatory property that unifies paradigm instances of gender or race? As I said, we can appeal to some central descriptions associated with our concepts so as to rule out any candidate referents that do not satisfy some of those central, hard to give up descriptions. But this could still leave room for disagreement regarding which properties can better explain what unifies those paradigm cases.

One proposal could be to appeal to the kind of metaphysical project that Ted Sider (2011) has recently developed (drawing on Lewis 1983). In a nutshell, his main idea is that when it comes to figuring out the meaning of a certain relevant term, we should always choose the most joint-carving candidate (among those candidate meanings that satisfy the inferential role associated with the term, if any). This view assumes that there are some descriptions of reality that are more joint-carving than others, where a concept or description is more joint-carving than another when it expresses a property or entity that is more fundamental than that expressed by the other. For instance, describing objects in terms of their being green or blue is more joint-carving than describing them in terms of their being *grue*. In this way, it could be argued that if we are wondering which objective type corresponds to our use of a concept C (as the descriptive project aims to reveal), we should find out which is the most joint-carving concept, out of the several candidate meanings under consideration.

⁵ The social constructionist has a possible alternative line of reasoning here: she might argue that there are no biologically significant properties in the vicinity of 'race' and 'gender' whatsoever, so that that the most objective properties in the vicinity can be social properties after all. This is a very controversial claim, and in my view social constructionists do not need to appeal to this possibility. As I have explained, another option would be to argue that our ordinary concepts of gender and race are associated with some central information that rules out any possible biological properties in the vicinity (if any) from being the referents.

For instance, we might wonder which is the most fundamental property, out of the several properties that are shared by paradigm instances falling under C. In my view, this proposal has some interesting commonalities with Haslanger's externalist account defended in her (2005, 2006). In the next section, I will develop this idea and I will explain how we can understand the notion of joint-carvingness in a way that can help us to make sense of the descriptive project of seeking the operative concept of gender and race. In section 5 I will argue that moral and political considerations can also be relevant with respect to the descriptive project, according to this framework.

3. Objectivity and Joint-Carvingness

Sider (2011) famously begins his book *Writing the Book of the World* by claiming that metaphysics is mainly concerned with the fundamental structure of the world. However, Elizabeth Barnes (2014) has rightly pointed out that disputes about the nature of gender and race are not disputes involving terms that are perfectly joint-carving. Indeed, many interesting philosophical disputes, such as debates about the nature of the mind, consciousness, concepts, meaning, species, composites, and artifacts, to name just a few, are arguably not about fundamental matters, and the corresponding terms are arguably not perfectly joint-carving. Therefore, if metaphysics is only concerned with the fundamental, we couldn't have proper metaphysical debates about gender and race (or many other non-fundamental phenomena). It seems plausible to conclude that we should revise our conception of metaphysics so as to allow debates about non-fundamental matters. However, the first paragraph of his book notwithstanding, Sider's characterization of metaphysics does not rule out the possibility of genuine metaphysical debates about non-fundamental matters. In this section, I will explain how his framework allows the possibility of debates about non-fundamental matters, and in particular I will argue that semantic externalists could understand the descriptive project of figuring out the referent of 'gender' and 'race' in terms of this framework.

So how can we have debates involving the notion of joint-carvingness, in order to talk about non-fundamental matters? Sider (2011), following Lewis (1983), appeals to the notion of terms that are not perfectly joint-carving but can still be *reasonably well* joint-carving, in the sense that they are more joint-carving than some alternatives.⁶ That is, when it comes to a certain level of inquiry, say, chemistry, or psychology, or sociology, we can always ask which terms are the most joint-carving with respect to that level. For many levels, all the terms that are posited by theories at that level are going to be less than perfectly joint-carving, but we can still compare them according to how well they carve nature at the joints. In this way, as Sider explains, the notion of fundamentality or joint-carvingness that he appeals to is *absolute* (i.e., an entity is either perfectly fundamental or not, and a term is either perfectly joint-carving or not), but at the same time we can introduce a comparative notion of less-than-perfect joint-carvingness, where two terms can be less than perfectly joint-carving but one

⁶ See Sider (2011: 77–8).

can be more joint-carving than the other (2011: 128–9). In what follows, I will say that a term that is not perfectly joint-carving but is more joint-carving than the (relevant) alternatives is *relatively* joint-carving. For example, terms that are relatively joint-carving with respect to sociology will arguably not be so with respect to chemistry, or physics. But we still can find out which terms are the most joint-carving at a certain non-fundamental level of inquiry, such as psychology or sociology. How can we figure this out? According to Sider (as I understand his view), this has to do with which terms or concepts are the most explanatorily useful, given the purposes and goals of explanations in that level. That is, to use Haslanger's terminology, when we are considering some terms in, say, sociology, and we are trying to figure out the operative concepts associated with them, a useful method here would be to figure out which kinds and properties (that are shared by the relevant paradigms) have the most explanatory power within sociology. In this way, we can make sense of the idea that in order to find out the operative concept, we need to find out which concepts are the most (relatively) joint-carving in that context, even if they are not perfectly joint-carving. So, of course, there are many debates in metaphysics that are not about perfectly joint-carving matters, but they can still be about relatively joint-carving matters.

One possible worry is the following: this framework appeals to different levels, such as physics, chemistry, biology, psychology, sociology, and so on, which are supposed to be ordered in a hierarchical way, according to how fundamental they are. And in addition, I am trying to make sense of the idea that there are some terms or concepts pertaining to a certain level that are more fundamental or more joint-carving than others with respect to that very same level. How can we make sense of this framework? It seems that we would need two different notions of fundamentality: on the one hand, we need a notion in order to account for the idea that physics is more fundamental than chemistry, which in turn is more fundamental than biology, and so on; and, on the other hand, we need a separate notion in order to make sense of the idea that we can compare different concepts posited at a certain level, according to how relatively fundamental or joint-carving they are, that is to say, according to how explanatorily useful they are. But it is not clear at all that we can make sense of this complex framework just in terms of the elusive notion of relatively fundamental or relatively joint-carving, in the sense of relative explanatory power I have suggested. For if we understand relative fundamentality in terms of different levels, then how can we compare the relative fundamentality of properties within the same level?

In my view, a tentative solution would be the following. We can use a version of the layered model of reality that Jaegwon Kim (2002) has explored in detail, following Oppenheim and Putnam (1958). According to this model, we can make sense of different layers or levels of reality in terms of the part-whole relation, that is, in terms of a mereological conception of reality. This model is primarily an account of different levels of reality in terms of different sums of entities. The model goes roughly as follows: at the most fundamental level n we have the most basic particles, which (we are assuming) cannot be decomposed any further, that is, they are not composed of further entities. At the following level, $n+1$, we have entities that are entirely decomposable into entities at level n . At level $n+2$, we have entities that are

entirely decomposable into entities at level n and $n+1$, and so on. For instance, at a certain level m (where $m > 2$) we will have molecules, which are decomposable into entities of lower levels. But the proper entities of level n and level $n+1$ are not decomposable into molecules. Likewise, at a certain level r (where $r > m$) we will have cells, which are decomposable into entities of the lower level m , but the proper entities of that lower level m are not decomposable into cells. This is what gives rise to the hierarchy of levels. As we have seen, this model makes use of the idea of a part-whole relation in order to account for the different levels of reality.⁷

Once we assume this mereological conception of levels, we can make sense of the idea of concepts that are more fundamental or more joint-carving than others at the same level of reality. First, we can talk about properties that appear only at certain higher levels, such as being a cell, or being a human being, or being conscious. We can then introduce new concepts in order to pick out these properties, and we can then talk about concepts that are posited at certain levels (when they refer to properties that appear in those levels). The crucial idea is that we can compare different concepts of a certain non-fundamental level (i.e., a level other than n) according to their explanatory power, that is, according to how explanatorily useful they are, with respect to the aims and purposes of our inquiries at that level.

Our next question is: what kinds of answers can we get when this method is applied? There are two main possible answers for a given descriptive project: either the term has a unique referent, that is, there is a candidate meaning that is the most objective or explanatorily useful, or the term is such that there are several candidate meanings that are equally joint-carving (with respect to that level). As Sider (2011) explains, candidate meanings are possible referents in the vicinity that satisfy the descriptions or inferential role associated with the concept, if any. If the paradigms have more than one property in common, where all of them satisfy the descriptions associated with the term (if any) and are equally objective or explanatorily useful, then we can say that there are several candidate meanings that are equally joint-carving. On the other hand, if there is a unique candidate meaning that is the most explanatory, then this gives rise in turn to two possibilities again: either this unique referent (the operative concept) corresponds to the referent the term should have (i.e., the target concept), or not. When the operative and the target concept come apart, this corresponds to the standard conception of the ameliorative project: a certain term actually refers to a certain property but given moral and political considerations, we ought to revise the meaning. In the next section I will argue that normative considerations are also relevant for the other outcomes, in addition to

⁷ As Kim (2002) argues, this mereological model of levels does not correspond to the standard division of levels in terms of physics, chemistry, biology, and so on. For it is not clear that biology only posits entities of a certain mereological order. It seems that the contrary is true: biology can talk about micro-physical entities such as bacteria, about medium-sized objects such as human beings, and about larger entities such as entire populations, for example species. For this reason, the mereological model of levels and the standard hierarchy of levels of study do not perfectly match. I am hoping this is not a big problem for my account here. What is crucial is that we have a way of distinguishing between different levels of inquiry vs. different degrees of explanatory power. We could appeal to the standard division of levels, in a somewhat unprincipled way, or we could appeal to this mereological model of levels, which is a bit more motivated, but less standard. Either way, I hope this does not affect the usefulness of my approach here.

this standard conception of the ameliorative project. In particular, I will argue (i) that in order to figure out the operative concept associated with a term in the first place, we need to appeal to normative considerations; and (ii) that in cases where a term has several candidate meanings that are equally joint-carving, and therefore (on Sider's view) the referent is indeterminate, we should also appeal to normative considerations in order to check whether one of the candidate meanings ought to be the referent rather than the others. Therefore, normative considerations are not only the business of ameliorative projects as traditionally conceived, but also the business of descriptive projects.

4. Reference, Explanatory Power, and Normative Considerations

As we have seen, according to Haslanger (2005, 2006), it could turn out that the operative concepts of gender and race correspond to social constructivist analyses of gender and race (as characterized in section 1). This will depend on whether these socially constructed properties are the most objective types unifying paradigms of genders and races, that is to say, on whether those socially constructed properties are the most explanatorily useful. But how can we settle this question? In her (2000) paper, Haslanger argued that those social constructivist analyses satisfy some important explanatory roles. Charlotte Witt (2011) has also argued that social constructivist accounts of gender can provide some crucial explanations of interest to feminist theories. However, Katharine Jenkins (2016) has recently argued that Haslanger's account of gender fails to do justice to some important explanatory requirements, namely, to capture the gender identity of trans women who self-identify as women and therefore as being the same gender as cisgendered women. Jenkins argues, very convincingly in my view, that an account of gender that fails to do justice to this requirement cannot be satisfactory. In particular, she argues that we need two notions of gender, namely, a notion of gender as a social class (in terms of Haslangerian social structures) and a notion of gender identity in terms of self-identification, in order to have a satisfactory account of gender. But a worry arises here: it is not clear whether we can capture this important insight purely in terms of the externalist framework that Haslanger (2005, 2006) put forward. As I suggested above, a good way of making sense of the externalist framework is to argue that our terms fix their referent in an externalistically determined way, that is, by means of causal relations to the most explanatorily useful kinds in the vicinity, and so on. (This is basically a form of *reference magnetism*: our terms latch on to the most objective kinds in the vicinity, where these are understood as the ones that satisfy certain explanatory roles, enter in law-like generalizations, and so on.⁸) And it is not obvious that these considerations will always be sufficient in order to identify a unique property corresponding to our notion of gender or race.

As I mentioned above, there are two cases that I want to consider. First, it is at least conceivable that the explanatory, empirical reasons of the sort advocated by Sider's

⁸ See Sider (2011: section 3.2), for further elaboration.

framework turn out to establish that our operative concept here corresponds to the social constructivist analysis proposed by Haslanger. This is, after all, the option suggested by Haslanger (2005, 2006), and as she already anticipated in her (2000), social constructivist analyses of gender do satisfy an important explanatory role. If that were the case, that is, if the objective type that our current usage of the concept 'gender' actually tracks corresponds to Haslangerian social structures, and if Jenkins (2016) is right when she argues that Haslanger's account of gender does not include all trans women under the corresponding group for 'woman', then our current usage of the term is arguably morally and politically objectionable. Many feminist theorists have argued (including Bettcher 2009; Saul 2012; Barnes 2014; Díaz-León 2016; and Jenkins 2016) that it is morally problematic to use a concept of 'woman' according to which there are some trans women that do not count as women (at least in most contexts).⁹ Therefore, it could be argued that if the operative concept associated with 'gender' corresponds to Haslangerian social structures (or some other kinds that exclude trans women), then we have strong moral and political reasons to change the meaning of our term 'gender' so as to make it more inclusive. This is the business of the ameliorative project, or to use the label recently coined by Burgess and Plunkett (2013a,b), the business of *conceptual ethics*.

But I also want to argue that normative considerations of the sort that are relevant in ameliorative projects or conceptual ethics more in general could be relevant with respect to the other semantic possibilities too. Let's go back to the original question: we were wondering what the operative concept of gender is, and whether it corresponds to Haslangerian social structures, or Jenkins' notion of gender identity, or something else.¹⁰ One possibility might be that when we appeal to the sort of explanatory reasons that Haslanger and Sider advocate, they just do not suffice to identify a unique kind as the most objective, explanatorily useful kind in the vicinity. For example, it might be that different kinds are useful for different purposes, and there is no clear way of choosing among those different purposes, given purely explanatory, empirical reasons. That is, in this case we have a term such that all the relevant candidate meanings are equally joint-carving. (As I said above, the relevant candidate meanings here are those properties that are shared by the paradigms and satisfy the descriptions or the inferential role associated with the term, if any.) What should we do in cases of indeterminacy of this sort? In my view, the best option is to appeal to the relevant moral and political reasons in the vicinity in order to decide which candidate meaning *should* be the meaning of that term (if any). This is part of the business of ameliorative projects or conceptual ethics, but it is different from the previous option in one crucial respect: instead of replacing the (determinate) meaning of a certain term with a different meaning, we are recommending that a

⁹ Giving a full defence of this claim is outside the scope of this chapter, and I take it to be obvious anyway. But some of the reasons are the following: claiming that trans women are not women helps to promote and perpetuate stigma and discrimination against trans women, which results in great harm, exclusion, and even violence. A notion of gender that fails to take this into account is contributing to this harm and therefore is morally problematic.

¹⁰ See Mikkola (2011), Sveinsdóttir (2011), and Witt (2011), for alternative accounts of gender; and Saul (2012), Bettcher (2013), and Díaz-León (2016) for alternative accounts of the meaning of 'woman'.

certain term with several candidate meanings (where it is indeterminate which candidate meaning it picks out) comes to have a unique determinate meaning, out of those alternatives.

We are now in a position to see that normative considerations of the sort that conceptual ethics emphasizes are relevant at the different stages of a metaphysical inquiry of the form ‘What is X?’, or ‘Is X real?’. In order to build my case for this claim, I will explain again what the different stages of a metaphysical project of this sort amount to, and what kind of normative considerations can be relevant for each stage.

First of all, we have the *conceptual* project of finding out what central, hard to give up information ordinary speakers associate with the term ‘X’. In my view, normative considerations such as pragmatic constraints, and even moral and political considerations, could be relevant at this earlier stage too. For example, in some cases we have concepts that are just associated to very thin descriptions such as ‘whatever turns out to be the most objective kind shared by those paradigms’. In cases like this, it will make a big difference which paradigms we choose. It could be argued that in cases of some contested terms such as ‘woman’ or ‘white’, we have several choices about which paradigms we should focus on, and there are moral and political considerations that are relevant here. For instance, Bettcher (2013) argues that there are different communities that use the term ‘woman’ in different ways, and in particular she identifies a dominant conception of women that excludes some trans women (which is perhaps a more widespread usage), and a resistant conception that includes all trans women (which is perhaps only taken up by a minority of speakers). It could be argued that there are moral and political considerations that would recommend focusing on the community of speakers that endorse the resistant conception instead of the dominant conception. If we focus on the resistant conception, then arguably our class of paradigms can include trans women. As Bettcher puts it, “it is inappropriate to dismiss alternative ways in which those terms are actually used in trans subcultures; such usage needs to be taken into consideration as part of the analysis” (2013: 235). These considerations concern how to understand the very descriptions that we associate with the term (e.g., “the most objective property that is shared by *these* paradigms”), so they belong to the *conceptual* stage of figuring out the application conditions of the concept. We can then conclude that moral and political considerations can be relevant with respect to the first step of the descriptive project, that is to say, the project of finding out the application conditions of the term.

Second, there is the *empirical* stage of figuring out what objective type actually satisfies those application conditions that we have established in the first stage. In order to find out what is the most objective type in the vicinity (among those that satisfy the application conditions), we will need to figure out what is the most explanatorily useful kind shared by the paradigms. In my view, in order to be able to compare different properties in virtue of how explanatorily useful they are, we need to decide first what are the main goals and purposes of our inquiry at that level. Only then can we figure out which explanations are more useful, *with respect to the relevant goals and purposes in that context*. In the remainder of this section I will explain my main argument for this claim.

As we saw above, in order to make sense of disputes about non-fundamental matters we have to find out which candidate meanings are more joint-carving, with respect to a certain non-fundamental level. But explanations at different non-fundamental levels could have different aims and purposes. First, different levels have different *explananda*, for instance, biology aims to explain the behavior of biological populations whereas sociology aims to explain the behavior of social entities such as social groups or institutions. Actually, it could be argued that these different explananda might concern entities at the same mereological level, say, groups of human populations, and if so, the only way of distinguishing the explanations would be in terms of the purposes of the explanation, rather than the level (which is the same). For instance, we could argue that what counts as explanatorily useful with respect to explaining and predicting the behavior of reproductively isolated biological populations does not correspond to what counts as explanatorily useful with respect to explaining and predicting the behavior of social groups (even if these happen to be exactly the same groups). That is, we can imagine that we are focusing on a certain class of properties, and we ask: which property is the most fundamental one, out of these? I want to argue that we cannot answer this question independently of the aims and purposes of our inquiry. If we are focusing on the aims and purposes of biological explanations, one of those properties might turn out to be the most explanatorily useful one, whereas if we are focusing on the aims and purposes of sociology, a different property (of the same set) might turn out to be the most explanatorily useful one with respect to that inquiry. This gives us reasons to deny Sider's claim that "the world has a distinguished structure, a privileged description. . . . There is an objectively correct way to 'write the book of the world'" (2011: vii).

Once we understand joint-carvingness and explanatory power in this pragmatic way, what considerations can be relevant in order to find out what concepts are the most explanatorily useful, with respect to a certain inquiry? In my view, both theoretical and pragmatic factors can be relevant here. We have to figure out which properties are the most explanatorily useful with respect to that inquiry, and we can make sense of this question only when we clarify what the aims and purposes of our explanations in that inquiry are. As Philip Kitcher (2007) aptly puts it: "there were lots of different ways in which the world of living things can be divided up, according to the things human beings find salient and according to the purposes they have" (p. 300).

David Ludwig (2015) has argued that there are *non-epistemic* values that are relevant with respect to the truth-value of scientific statements such as (1): 'There are two different tiger species in the San Diego zoo'. Here I want to extend his argument, in order to argue that there are non-epistemic values that are relevant with respect to the descriptive project of finding the referent of terms such as 'gender' and 'race'. Ludwig's argument has two main premises: (i) the truth-value of many scientific statements depends on which ontological framework we choose; and (ii) choices regarding ontological framework depend in part on non-epistemic values. Therefore, it follows that the truth-value of many scientific statements depends on non-epistemic values. For instance, the truth-value of (1) above depends

on which concept of species we choose, and this depends in turn on the explanatory interests of scientists.¹¹

I want to extend Ludwig's line of argument in order to apply it to our account of descriptive projects in philosophy. On my view, questions about what is the referent of a term will depend, first, on the application conditions for the term that we choose, and, second, on what turns out to be the most explanatorily useful property in the vicinity, where whether a property is more explanatorily useful than another depends in part on the explanatory interests of the inquirers. In this way, we can argue that the answer to descriptive projects in philosophy of the form 'Is X real?' (or what is equivalent, as Thomasson (2008) argues, 'Does 'X' refer to anything in the actual world?'), and 'What is X?' (or what I am assuming to be equivalent, 'What is the referent of 'X'?'), depends in part on non-epistemic values, including the explanatory interests that are relevant and appropriate for each project. Furthermore, we can argue that it is possible to assess the aptness of different explanatory goals according to different considerations, including theoretical, pragmatic, moral, and political factors (following Kitcher 2001). That is to say, moral and political factors can also be relevant at this stage, in two different ways: (a) in order to decide what are the most relevant aims and goals of our inquiry (i.e., some aims can be more politically useful than others); and (b) which explanations are the most useful with respect to some given goals (i.e., some explanations might be useful with respect to some criteria but not others, where these criteria are also morally and politically assessable). Therefore, we can conclude that normative considerations, including moral and political considerations, can be relevant with respect to a descriptive project in philosophy of the form 'What is X?' or 'Is X real?'. In particular, they are relevant in order to fix the criteria of explanatorily usefulness that we apply for comparing the different candidate meanings regarding how explanatorily useful they are.

In order to illustrate this line of argument, I will consider one example. We can consider the concept of *gender*, as the term was introduced in feminist theory in order to capture the sex/gender distinction. (See Mikkola 2008 for a historical survey.) Here I want to focus on the question: what is the operative concept associated with 'gender' in the context of feminist theory? I want to argue that in order to figure out the operative concept of 'gender', we need to make explicit the goals and purposes of explanations within feminist theory. Arguably, one of the main aims of feminist theory is to describe, explain, and resist the oppression of women. In my view, these aims can be used in order to compare different proposals about what is the most objective type that is shared by paradigms of 'gender'. This is one way in which normative considerations (including the sort of moral and political factors that motivate feminist theory) can be relevant in a descriptive project. We can compare this project with the different (descriptive) project of figuring out a notion of biological sex in biology. Arguably, one of the main aims of this inquiry is to explain sexual reproduction (although they might be others). In my view, the chosen goals will make a difference regarding what are the central explanatory aims and purposes

¹¹ Ludwig's excellent discussion draws on the arguments of many philosophers of science, including Dupré (1993); Anderson (1995); and Kitcher (2001).

with respect to which we should assess the explanatory usefulness of different candidate meanings for ‘male’ and ‘female’. It could also be argued that the paradigms that we should pick up are different in biology than in the case of the term ‘gender’ within feminist theory, and this would suffice to yield different referents. But even if we focus on similar paradigms, the shared properties that turn out to be the most objective ones with respect to ‘gender’ and ‘sex’, that is, the most explanatorily useful ones, could turn out to be different in each case, since the main aims and goals of our explanations in each area are different, and therefore what counts as explanatorily useful can be different.¹² But, to emphasize, these projects are descriptive projects if anything is, and therefore my point is that normative considerations are not only relevant with regards to ameliorative projects about what ‘gender’ should mean, but also regarding descriptive projects about what ‘gender’ actually means (in the context of feminist theory, say).

To recap: As I said above, there are two possible outcomes of the descriptive project: either there is a unique candidate meaning that is the most joint-carving (and therefore this will turn out to be the referent, due to reference magnetism), or there are several candidates that are equally explanatory, given our theoretical aims and constraints, and then it will be indeterminate what the referent is. In this second case, if there are any normative considerations that might be relevant here, including prudential, moral, and political reasons, then this would give us good reasons to choose one referent over the others, out of those equally objective candidate meanings, and in my view this means that we ought to revise the meaning so that the term gets to have that unique referent, instead of an indeterminate meaning.¹³

Then, we should reformulate the nature of descriptive projects seeking the operative concept as follows. The corresponding overall question will be the following: what is the most objective type (that satisfies the application conditions), that is, what is the most explanatorily useful kind, given all the relevant considerations, including theoretical, instrumental, moral, and political considerations? More generally, the aim here is to find out what is the most explanatorily useful kind in the vicinity, given what the relevant aims and purposes of our inquiry are, what the relevant paradigms are, and so on. And if it turns out that there are several candidate meanings satisfying

¹² Haslanger (2016) makes a similar case regarding the notion of sex. One of my aims here is to show that this line of argument can be generalized in order to provide a characterization of descriptive projects in philosophy that allows room for normative considerations, in a way that is compatible both with a radical externalist framework like Haslanger’s, and with a two-step descriptive project of the sort I have proposed.

¹³ There are at least two ways in which we could understand this indeterminacy. It could be the case that the referent of the term is genuinely indeterminate, in the sense in which vague terms such as ‘bald’ can have indeterminate referents. It could be argued that there are some possible precisifications of ‘bald’ that are more politically useful than others, and therefore this would give us some reasons for changing the meaning of ‘bald’ so that it comes to determinately refer to this precisification. Alternatively, we could say that the term is context-sensitive, such as ‘tall’. Arguably, ‘tall’ means something like ‘being taller than a certain threshold that is salient in this context’. In this case, it is not the case that the term has several candidate meanings such that it is indeterminate what the referent is at any given context. Rather, the term has a unique referent at each context where a threshold is made salient, but this threshold can change from context to context, and therefore the term can have different referents at different contexts. In Díaz-León (2016), I argued that moral and political considerations are relevant in order to determine which standards of similarity are more salient in each context, with respect to context-shifting terms.

the descriptions or information associated with the term that are all equally joint-carving (given the theoretical and practical criteria governing our inquiry), then we should appeal to any additional relevant moral and political considerations to check whether any of the candidate meanings can better satisfy these normative criteria.

Finally, it might be the case that even if there is a clear objective kind that is the most explanatorily useful corresponding to some term, according to the descriptive project, there might be additional moral and political considerations that trump this, so that it is the case that we *ought* to change the meaning of that term, all things considered. This is the most familiar version of an ameliorative project. But as I have argued, this is not the only stage where normative considerations are relevant. As we have seen, normative considerations are also relevant in order to find out what the actual referent of a certain term is, or what the referent should be, given a situation of indeterminacy.

5. Conclusion

We can conclude that, at the end of the day, there is no sharp distinction between debates that are properly descriptive, and debates that are ameliorative, since normative considerations are relevant at many different stages of both projects. In my view, it is more useful to see the distinction between descriptive and ameliorative projects as different stages of an overarching project, where the relevant overall question is the following: what is the most useful way of using ‘X’, or what should ‘X’ mean, given all relevant theoretical, practical, moral, and political considerations?

Acknowledgements

Research on this project has been funded by grants RYC-2012-10900 and FFI2015-66372-P (Spanish Government) and H2020-MSCA-ITN-2015-675415 (European Union). I have presented ancestors of this material at the Universities of Barcelona, Essen, Leeds, Leuven, Oslo, Sheffield and St Andrews, and the Eastern APA in Baltimore. I am very grateful to the audiences on all those occasions for very useful feedback. Extra thanks are due to Brice Bantegnie, Delia Belleri, Ingo Brigandt, Alexis Burgess, Herman Cappelen, Aurélien Darbellay, Sally Haslanger, Dan López de Sa, David Ludwig, David Plunkett, Kevin Scharp, Tim Sundell, and Amie Thomasson for very helpful comments.

References

- Anderson, E. 1995. Knowledge, Human Interests, and Objectivity in Feminist Epistemology. *Philosophical Topics* 23 (2):27–58.
- Barnes, E. 2014. Going Beyond the Fundamental: Feminism in Contemporary Metaphysics. *Proceedings of the Aristotelian Society* 104 (3):335–51.
- Bettcher, T. 2009. Trans Identities and First-Person Authority. In L. Shrage (ed.), *You’ve Changed: Sex Reassignment and Personal Identity* (pp. 98–120). Oxford: Oxford University Press.
- Bettcher, T. 2013. Trans Women and the Meaning of “Woman”. In N. Power, R. Halwani, and A. Soble (eds.), *The Philosophy of Sex* (pp. 233–49). London: Rowman & Littlefield.
- Burgess, A., and Plunkett, D. 2013a. Conceptual Ethics I. *Philosophy Compass* 8 (12):1091–101.

- Burgess, A., and Plunkett, D. 2013b. Conceptual Ethics II. *Philosophy Compass* 8 (12):1102–10.
- Chalmers, D., and Jackson, F. 2001. Conceptual Analysis and Reductive Explanation. *The Philosophical Review* 110 (3):315–61.
- Díaz-León, E. 2012. Social Kinds, Conceptual Analysis, and the Operative Concept: A Reply to Haslanger. *Humana.Mente: Journal of Philosophical Studies* 22:57–74.
- Díaz-León, E. 2016. *Woman* as a Politically Significant Term: A Solution to the Puzzle. *Hypatia* 31 (2):245–58.
- Dupré, J. 1993. *The Disorder of Things*. Oxford: Oxford University Press.
- Glasgow, J. 2009. *A Theory of Race*. London: Routledge.
- Haslanger, S. 2000. Gender and Race: (What) Are They? (What) Do We Want Them to Be? *Nous* 34 (1):31–55.
- Haslanger, S. 2003. Future Genders? Future Races? *Philosophic Exchange* 34 (1):4–27.
- Haslanger, S. 2005. What Are We Talking About? The Semantics and Politics of Social Kinds. *Hypatia* 20 (4):10–26.
- Haslanger, S. 2006. What Good Are Our Intuitions? Philosophical Analysis and Social Kinds. *Proceedings of the Aristotelian Society*, Sup. Vol. 80 (1):89–118.
- Haslanger, S. 2016. Theorizing with a Purpose: The Many Kinds of Sex. In C. Kendig (ed.), *Natural Kinds and Classification in Scientific Practice*. London: Routledge.
- Jackson, F. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.
- Jenkins, K. 2016. Amelioration and Inclusion: Gender Identity and the Concept of *Woman*. *Ethics* 126 (2):394–421.
- Kim, J. 2002. The Layered Model: Metaphysical Considerations. *Philosophical Explorations* 5 (1):2–20.
- Kitcher, P. 2001. *Science, Truth, and Democracy*. Oxford: Oxford University Press.
- Kitcher, P. 2007. Does 'Race' Have a Future? *Philosophy & Public Affairs* 35 (4):293–317.
- Lewis, D. 1983. New Work For a Theory of Universals. *Australasian Journal of Philosophy* 61:343–77.
- Ludwig, D. 2015. Ontological Choices and Value-Free Ideal. *Erkenntnis*. DOI: 10.1007/s10670-015-9793-3.
- Mikkola, M. 2008. Feminist Perspectives on Sex and Gender. In E.N. Zalta (ed.), *Stanford Encyclopaedia of Philosophy* (summer edn). <http://plato.stanford.edu/archives/sum2008/entries/feminism-gender/>.
- Mikkola, M. 2011. Ontological Commitments, Sex and Gender. In C. Witt (ed.), *Feminist Metaphysics* (pp. 67–83). New York: Springer.
- Mikkola, M. 2015. Doing Ontology and Doing Justice: What Feminist Philosophy Can Teach Us About Meta-Metaphysics. *Inquiry* 58:780–805.
- Oppenheim, P., and Putnam, H. 1958. Unity of Science as a Working Hypothesis. *Minnesota Studies in the Philosophy of Science*, vol. 2. Minneapolis: University of Minnesota Press.
- Saul, J. 2012. Politically Significant Terms and Philosophy of Language: Methodological Issues. In S. Crasnow and A. Superson (eds.), *Out from the Shadows: Analytical Feminist Contributions to Traditional Philosophy* (pp. 195–216). Oxford: Oxford University Press.
- Sider, T. 2011. *Writing the Book of the World*. Oxford: Oxford University Press.
- Sveinsdóttir, A. K. 2011. The Metaphysics of Sex and Gender. In C. Witt (ed.), *Feminist Metaphysics* (pp. 47–65). New York: Springer.
- Thomasson, A. 2007. *Ordinary Objects*. Oxford: Oxford University Press.
- Thomasson, A. 2008. Existence Questions. *Philosophical Studies* 141:63–78.
- Witt, C. 2011. *The Metaphysics of Gender*. Oxford: Oxford University Press.