

What Are We Talking About? The Semantics and Politics of Social Kinds

SALLY HASLANGER

Theorists analyzing the concepts of race and gender disagree over whether the terms refer to natural kinds, social kinds, or nothing at all. The question arises: what do we mean by the terms? It is usually assumed that ordinary intuitions of native speakers are definitive. However, I argue that contemporary semantic externalism can usefully combine with insights from Foucauldian genealogy to challenge mainstream methods of analysis and lend credibility to social constructionist projects.

When we talk of gender and race, at one level it is pretty clear what we're talking about. Although there are cases where it is hard to tell from casual observation what race or gender a person is, and although there are borderline cases in which our ordinary criteria don't give us a clear answer, we are all pretty well versed in the practice of assigning people a race and a gender. Yet, at another level, it is not so clear what we mean when we say "I'm a white woman" or "Barack Obama is a black man." For example, race eliminativists maintain that talk of races is vacuous (no one is white or black, Asian or Latino, because there are no races); others argue that race continues to be a meaningful biological kind; and still others argue that race is a social category. Feminists have questioned the legitimacy of dividing us into two sexes, males and females, and many have grown dubious of the sex/gender distinction altogether; in everyday discourse the term 'gender' now seems to be equivalent to 'sex'; and yet many feminist theorists still argue that gender is a social category. How do we make sense of all this? Are the apparent disagreements real disagreements, or are the different parties to these discussions really talking about different things?

Elsewhere I've defended social constructionist accounts of race and gender (Haslinger 2000). I believe that races and genders are real categories to be

defined in terms of social positions. I have come to this conclusion by considering what categories we should employ in the quest for social justice. Although I believe there is reason to conclude that biological essentialism about race and gender is false, to deny that people are raced and gendered within (at least) the contemporary United States would be to ignore facts about our social arrangements that those who seek justice cannot ignore. On my view, to say that I am a white woman is to situate me in complicated and interconnected systems of privilege and subordination that are triggered by interpretations of my physical capacities and appearance. Justice requires that we undermine these systems, and in order to do so, we need conceptual categories that enable us to describe them and their effects. A consequence of my view is that when justice is achieved, there will no longer be white women (there will no longer be men or women, whites or members of any other race). At that point, we—or more realistically, our descendants—won't need the concepts of race and gender to describe our current situation. However, we (they) will probably need the concepts in order to understand our past, just as, for example, to make sense of American social history, it is valuable to have the concept of 'quadroon,' 'octoroon,' 'spinster,' and the like.

Much recent debate over race, in particular, seems to have become bogged down in the question whether this or that account of race can claim to be an analysis of *our* concept of race (See, for example, Mallon 2004, Hardimon 2003). In developing constructionist accounts of race and gender, I've maintained that my goal is *not* to capture the ordinary meanings of 'race' or 'man' or 'woman', nor is it to capture our ordinary race and gender concepts. I've cast my inquiry as an analytical—or what I here call an *ameliorative*—project that seeks to identify what legitimate purposes we might have (if any) in categorizing people on the basis of race or gender, and to develop concepts that would help us achieve these ends. I believe that we should adopt a constructionist account not because it provides an analysis of our ordinary discourse, but because it offers numerous political and theoretical advantages.

However, in this essay, I want to reconsider the strategy behind my own proposals, and social constructionist proposals more generally, and argue that they stand in a more complicated relationship to the project of analyzing ordinary discourse or explicating our concepts than I previously suggested. In doing so, I will offer a framework that clarifies the relationship between social constructionism and other philosophical projects, both naturalistic and a priori. The broad goal of this paper is to question what's at issue in doing philosophical analysis of a concept, and to disrupt the assumptions behind the common revisionary/nonrevisionary contrast.

I begin by sketching a number of different projects that might legitimately count as providing an analysis of our concepts or speech. It is by now a familiar theme in philosophy of language that meanings (or at least some meanings)

“aren’t in the head”; yet it is a complicated matter to figure out the relationship between what is in our heads and the content of what we say, and think, and do. When thinking about socially and politically meaningful concepts, we must also be attentive to the possibility that what’s in our heads may not only be incomplete, but may be actively masking what’s semantically going on. Part of the job of ideology may be (somewhat paradoxically) to mislead us about the content of our own thoughts. How can we make sense of this? And, if this is the case, what becomes of the project of philosophical analysis?

GENEALOGY: TARDINESS

The project of conceptual analysis in philosophy takes many forms, partly depending on the particular concept in question, and partly depending on what methodological assumptions the philosopher brings to the issue. There are at least three common ways to answer “What is X?” questions: *conceptual*, *descriptive*, and *ameliorative*.¹

For example, consider the question: What is knowledge? Following a *conceptual* approach, one is asking: What is *our* concept of knowledge? and looks to *a priori* methods such as introspection for an answer. Taking into account intuitions about cases and principles, one hopes eventually to reach a reflective equilibrium. On a *descriptive* approach, one is concerned with what kinds (if any) our epistemic vocabulary tracks. The task is to develop potentially more accurate concepts through careful consideration of the phenomena, usually relying on empirical or quasi-empirical methods. Scientific essentialists and naturalizers, more generally, start by identifying paradigm cases—these may function to fix the referent of the term—and then draw on empirical (or quasi-empirical) research to explicate the relevant kind to which the paradigms belong. Paradigms for knowledge could include my knowledge that there is a pencil on the desk in front of me, my daughter’s knowledge that $2 + 2 = 4$, the scientist’s knowledge that $E = mc^2$, a sampling of further cases of memory, testimony, and the like. The question is whether these states form a natural kind, and if so, what kind? A descriptive approach in philosophy of mind and epistemology sometimes draws on cognitive science.

Ameliorative projects, in contrast, begin by asking: What is the point of having the concept in question—for example, why do we have a concept of knowledge or a concept of belief? What concept (if any) would do the work best? In the limit case, a theoretical concept is introduced by stipulating the meaning of a new term, and its content is determined entirely by the role it plays in the theory. If we allow that our everyday vocabularies serve both cognitive and practical purposes that might be well-served by our theorizing, then those pursuing an ameliorative approach might reasonably represent themselves as providing an account of our concept—or perhaps the concept we are reaching for—by enhancing our

conceptual resources to serve our (critically examined) purposes. Conceptual, descriptive, and ameliorative projects cannot, of course, be kept entirely distinct, but they have different subject matters and different goals.

In this essay, I consider an additional approach: *genealogy*. Later, I consider whether it should be considered a more specific form of the three approaches just mentioned, or in a distinct category. The idea of a genealogical approach stems from Nietzsche and Foucault, though it has been taken up by a wide range of scholars in the humanities. Very roughly, a genealogy of a concept explores its history, not in order to determine its true meaning by reference to origins, and not for sheer historicist fascination, but in order to understand how the concept is embedded in evolving social practices. Two points are crucial here: First, our concepts and our social practices are deeply intertwined. Concepts not only enable us to describe but also help structure social practices, and our evolving practices affect our concepts. Second, there is often a significant gap between the dominant or institutional understanding of a domain and its actual workings, for example, in the interplay between concept and practice, developments on one side can get ahead of or stubbornly resist the other.

For example, in some school districts, there are complex rules and consequences constructed around the notion of being *tardy*. There are forms for tracking tardiness; school officials looking out for tardiness; if you are tardy too many times in a year, you can be suspended or expelled, can't be promoted to the next grade, and so on. In school districts where this is the case, there are local understandings of how to navigate the system. For example, one morning when we were running especially late, my son Isaac reassured me by saying, "Don't worry Mom, no one is ever tardy on Wednesdays because my teacher doesn't turn in the attendance sheet on Wednesday until after the first period." This fact, together with the knowledge that his teacher would mark him present as long as he arrived before the attendance sheet was turned in, meant that *in practice* 'tardy' was defined differently in his classroom from the way it was, say, in the classroom next door.

How should we understand this? It might be tempting to insist that Isaac *really was* tardy when he arrived after the bell, even if his teacher didn't mark him as such on the attendance sheet. In other words, there is one *real* definition of tardy (the school district's: any student arriving in his or her homeroom after the 8:25 A.M. bell is tardy), and the others are only approximations and would be recognized as such by those involved.² However, we should note that such insistence would involve privileging the explicit institutional definition of *tardy* over the more implicit meaning established within the particular classroom practice.³ In a slightly different context, one might imagine a teacher arguing with an overzealous school official by saying something like: "Yes, Sophia arrived two minutes after the bell rang, but students were still hanging up their coats. She wasn't *tardy*."

A genealogical approach is interested in the social and historical circumstances that give rise both to the disciplinary structures within which tardiness has its institutional meanings and to those that give rise to alternative, sometimes subversive, practices that arise in the day-to-day lives of those within the institution.⁴ So in a genealogical account of ‘tardiness’ one would expect to find a story about how various conceptions of ‘tardy’ are embedded in the evolution of multiple and interacting social practices. My point is not to argue that either the classroom or the school district definition should be privileged; rather, (at the moment) it is to highlight that tardiness plays a role in different, and in some cases competing, practices.

In the literature on genealogy, the relevant contrast is often taken to be between broad institutional meanings and alternative local ones. However, this is one of several different axes of comparison that might be relevant. For example, in general, when we consider the use of terms or concepts in context there are important differences between:

- institutional uses v. “local” uses
- public uses v. more idiosyncratic individual uses
- what is explicit v. what is implicit in the minds of users
- what is thought (what we take ourselves to be doing with the concept) v. what is practiced (what we’re actually doing with it)
- appropriate v. inappropriate uses

In the case of ‘tardy,’ the school board’s notion is public, explicit, more often recited than practiced and, one might think, an overly rigid definition of what tardiness really is (recall the teacher’s complaint on Sophia’s behalf); the local classroom notion is less public (though not private), implicit, more often practiced than recited, and, one might think, an overly ad hoc understanding of what tardiness really is (you’re tardy unless you arrive around 8:25 on Mondays, Tuesdays, Thursdays, and Fridays, and before 9:00 A.M. on Wednesdays). Although a concern with power may recommend being especially attentive to the distinction between institutional and local meanings, for our purposes it will be important to have available the distinction between what I’ve elsewhere called the *manifest* concept and the *operative* concept (Haslanger 1995). Roughly, the manifest concept is the more explicit, public, and “intuitive” one; the operative concept is the more implicit, hidden, and yet practiced one.⁵

Although I’ve focused on the simple example of ‘tardy,’ there are, of course, more philosophically rich examples available. Feminist and race theorists have been urging for some time that the proper target of analysis is not (or not simply) what we have in mind, but the social matrix where our concepts do their work. For example, Catherine MacKinnon says the verb *to be* in feminist

theory “is a very empirical ‘is.’ Men define women as sexual beings; feminism comprehends that femininity ‘is’ sexual. Men see rape as intercourse; feminists say much intercourse ‘is’ rape” (MacKinnon 1987, 59). Charles Mills argues that the Enlightenment social contract is a racial contract (Mills 1997), and that an adequate analysis of personhood reveals that “all persons are equal, but only white males are persons” (Mills 1998, 70). Such analyses purport to show that our manifest understandings of crucial political notions are masking how the concepts in question actually operate (see also Mills 1998, 139–66).

It is important to note, however, that the axes of comparison I’ve listed introduce a contrast between what tardiness (femininity, personhood) “really is” and the competing understandings of tardiness used in practice that takes us beyond genealogy. Within a genealogical inquiry our subject matter is a set of historically specific social practices. To give an account of *what tardiness really is*, is to describe a broad matrix of practices, procedures, rules, rationales, punishments, institutions, equipment (bells, clipboards, forms), to demonstrate how power circulates within it, and how certain subject positions (the walkers, the bus-riders, the habitually tardy) are formed (see also Hacking 1999, 10–14). On the genealogical approach, this matrix is what tardiness *really is*.⁶

However, in suggesting above that both the local and institutional definitions of tardiness were in some respects inadequate, I was implying that there is a further way of thinking about what tardiness “really is” that should take us into normative questions: Should we have the category of ‘tardy’ in our school district? If so, how should it be defined? One might be tempted to think that the situation in our local school is ripe for an ameliorative inquiry that would have us consider what the point is of a practice of marking students tardy, and what definition (and corresponding policy) would best achieve the legitimate purposes.

The lack of attention to the normative is the basis for an important and influential criticism of genealogical inquiry. Although genealogy is attentive to and describes the use of normative discourse and the impact of social norms, it attempts to forswear making normative claims; as a result, it cannot make crucial distinctions between good and bad forms of power and authority, legitimate or illegitimate force (Fraser 1989, chap. 1). Correlatively, one might complain that analytic inquiry that attempts to improve on our current definitions typically fails to understand how our current concepts have structured our practices, distribute power and authority, and bring with them false assumptions of legitimacy. It is tempting to think that genealogy without normative analysis shirks its responsibilities; and normative analysis without genealogy is out of touch with reality. Note again that to distinguish the variety of philosophical projects is not to say that they can or should be pursued independently; yet making clear the differences can help us locate our disagreements and misunderstandings.

FORMS OF GENEALOGY

So, what is the relationship between genealogy and the approaches mentioned earlier, namely, the conceptual, descriptive, and ameliorative? Insofar as the goal of genealogy is to understand how concepts are embedded within social matrices, it is possible to modify any of these more traditional approaches in the spirit of genealogy.

For example, the conceptual approach I've described focuses on a priori reflection and ideas that are relatively accessible to introspection; it is plausible to see this as an investigation of the manifest concept.⁷ In undertaking conceptual analysis of, say, *F*-ness, it is typically assumed that it is enough to ask competent users of English under what conditions someone is *F*, without making any special effort to consult those whose daily lives are affected by the concept. However, if one is sensitive to the possibility that in any actual circumstance there are competing meanings (often quite explicit) that structure alternative practices, then it seems worth considering a broad range of speakers, who are differently situated with respect to the phenomenon. A *conceptual genealogy* of 'tardy' would not be content with reflection by a competent English speaker, but would require attention to differently situated speakers over time. We would need to ask: What are the range of meanings? Whose meanings are dominant and why?

Of course, some speakers may not be very thoughtful about their use of terms, and others may simply be confused. Yet we should keep in mind that "our" concept may not be univocal; in our haste to find a univocal concept, we may obscure how the concept works in a complex social context. Such investigations into a broader range of ideas and practices will not only be relevant to a conceptual genealogy, but also to an *ameliorative genealogy* that undertakes to evaluate the point of having a concept or structure of concepts (along with related practices) and proposes improved resources to fulfill them.

In this essay, however, I am especially interested in exploring how genealogy might affect a descriptive approach. Those pursuing a descriptive approach will usually select paradigms from commonly and publicly recognized cases; as suggested before, the task is to determine the more general type or kind to which they belong. For example, the case in which Isaac arrives at school at 8:40 A.M. (when school starts at 8:25 A.M.) would count as a paradigm case of tardiness, regardless of what his teacher marks on the attendance sheet. Of course, the aim of a descriptive project in this case is not to provide a *naturalistic* account of tardiness—one that would seek to discover the *natural* (as contrasted with *social*) kind within which the paradigms fall—given that the notion of being "on time" concerns one's behavior in response to a complex set of norms and expectations. But it is possible to pursue a descriptive approach within a social domain as long as one allows that there are social kinds or types.⁸ In fact, I've chosen to speak of descriptive approaches rather than naturalistic ones for just this reason. Descriptive analyses of social terms such as 'democracy' and

'genocide' or ethical terms such as 'responsibility' and 'autonomy' are methodologically parallel to more familiar naturalizing projects in epistemology and philosophy of mind. However, the investigation of social kinds will need to draw on empirical social/historical inquiry, not just natural science.

If one were to undertake a *descriptive genealogy* of 'tardiness,' then it makes most sense to start with a social context in which tardiness plays a role. The first task is to collect cases that emerge in different (and perhaps competing) practices; then, as before, one should consider if the cases constitute a genuine type, and if so, what unifies the type. This, of course, cannot be done in a mechanical way and may require sophisticated social theory both to select the paradigms and to analyze their commonality; and it is easily possible that the analysis of the type is highly surprising. For example, it was not intuitively obvious that water is H_2O or that gold is an element with the atomic number 79. It took sophisticated natural science to determine what the terms 'water' and 'gold' mean. In any descriptive project, intuitions about the conditions for applying the concept should be considered secondary to what the cases in fact have in common; so as we learn more about the paradigms, we learn more about our concepts.

SEMANTIC EXTERNALISM

I've suggested that there are different projects that might count as attempting to theorize what tardiness is. Because these projects will reasonably yield different accounts, one might wonder which strategy is entitled to claim that its results provide an analysis of the concept. The problem should look more familiar if we situate this discussion in the tradition of semantic externalism. Externalists maintain that the content of what we think and mean is determined not simply by intrinsic facts about us but at least in part by facts about our environment. Remember: Sally and Twinsally both use the term 'water,' but Sally means H_2O and Twinsally means XYZ (Putnam 1975). Sally thinks she has arthritis in her thigh, and is wrong because 'arthritis' in her environment is an ailment of the joints; Twinsally thinks she has arthritis in her thigh and is right because 'arthritis' in her environment is an ailment that is not confined to the joints (Burge 1979).

Externalism initially appeared in two forms, supported by the sorts of examples just recited:

Natural kind externalism (Putnam 1975; Kripke 1980): natural kind terms or concepts pick out a natural kind, whether or not we can state the essence of the kind, by virtue of the fact that their meaning is determined by ostension of a paradigm (or other means of reference fixing) together with an implicit extension to "things of the same kind" as the paradigm.

Social externalism (Putnam 1973; Burge 1979): the meaning of a term or the content of a concept used by a speaker is determined at least in part by the standard linguistic usage in his or her community.

It then became clear that externalist phenomena are not confined to natural kind terms (properly speaking) but occur quite broadly. For example, in the history of logic and math, inquiry can seem to converge on an idea or concept that we seemed to have in mind all along, even though no one, even the best minds, could have explicated it. (Leibniz's early efforts to define the limit of a series is an example.) In such cases, it is plausible to maintain that certain experts were "grasping a definite sense, whilst also failing to grasp it 'sharply'" (Peacocke 1998, 50). Although Fregeans are apt to capture this by invoking objective senses that the inquirers "grasp," an ontology of sparse objective properties will also do the work.

The upshot of this is that the basic strategy of natural kind externalism need not be confined to natural kinds (where it is assumed that things of the same natural kind share an essence). Externalism is an option whenever there are relatively objective types. The notion of objective type needed is not too mysterious: a set of objects is more an objective type by virtue of the degree of unity amongst its members beyond a random or gerrymandered set. We might account for unity in various ways (Lewis 1983), but a familiar way I'll assume for current purposes is in terms of degrees of similarity; the similarity in question need not be a matter of intrinsic similarity, that is, things can be similar by virtue of the relations (perhaps to us) they stand in. Roughly,

Objective type externalism: terms or 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.

Sets of paradigms 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. We should not assume that objectivity is only found in the natural world. There are objective types in every realm: social, psychological, political, mathematical, artistic, and so on.⁹

What does externalism have to do with genealogy? Genealogy explores the embeddedness of a concept within social practices and the history of those practices. Just above I suggested that a *conceptual genealogy* would explore the relatively explicit ideas and assumptions associated with a concept (over time),

taking into account how these may vary depending on one's position within the practice structured by the concept. A *descriptive genealogy* explores how a term functions in our evolving practices and manages to pick things out. Descriptive projects, of the sort I've indicated, adopt an externalist approach to content: they set out to determine the objective type, if any, into which the paradigms of a particular concept fall. Descriptive projects become genealogical to the extent that they attend to the concrete historical workings of our practices and how the concept is actually used to structure our ongoing activities. In effect, a descriptive project will aim to disclose the *operative* concept(s), while the conceptual project explicates the manifest.

In some cases, the manifest concept and operative concept coincide: when we are clear what exactly we are talking about. But in many cases a speaker could have as the content of her thought or speech something about which she was ignorant or even seriously misguided. Given the externalist backdrop, this is not surprising. As the externalist slogan goes, "Meanings ain't in the head." The genealogist is especially keen to explore cases in which the manifest and operative concepts come apart, that is, when the operation of the concepts in our lives is not manifest to us. If one assumes that the task of philosophical inquiry is simply to explicate the dominant manifest meaning of a term, then any genealogical inquiry—almost any externalist inquiry—will seem revisionary. But philosophical inquiry—even philosophical inquiry that takes its goal to be the analysis of our concepts—should not define itself so narrowly, or else it is in danger of collapsing into lexicography (an interesting endeavor, to be sure, but not our only option).

DESCRIPTIVE GENEALOGIES OF RACE AND GENDER

I've suggested so far that there are several different projects that might plausibly be thought to provide an analysis of our concepts, and several different kinds of subject matter that might be analyzed.¹⁰

Conceptual analyses elucidate "our" (manifest) concept of *F*-ness by exploring what "we" take *F*-ness to be.

Conceptual genealogy: elucidate the variety of understandings and uses of *F*-ness over time and across individuals differently positioned with respect to practices that employ the notion.

Descriptive analyses elucidate the empirical kinds (the operative concept) into which "our" paradigm cases of *F*-ness fall.

Descriptive naturalism: elucidate, where possible, the *natural* (chemical, biological, neurological) kinds that capture "our" paradigm cases of *F*-ness.

Descriptive genealogy: elucidate the *social matrix* (history, practices, power relations) within which “we” discriminate between things that are *F* and those that aren’t.

Ameliorative analyses elucidate “our” legitimate purposes and what concept of *F*-ness (if any) would serve them best (the target concept). Normative input is needed.

Although I have distinguished the different projects and subject matter, there will be cases in which they completely coincide. In other words, there will be cases in which we are aware of what we are talking about, and what we are talking about is what we should be talking about, namely, where the manifest, operative, and target concepts are the same. There will be cases in which an ameliorative project targets the kind that we are, and take ourselves to be, tracking. But there will also be times when these come apart, for example, where ignorance or ideology masks what we are doing or saying.

When the manifest, operative, and target concepts come apart, there will be different ways to unite them. For example, if the target concept and manifest concept coincide and it is our practice that fails, the best strategy is plausibly to correct the practice to meet the standards we ourselves affirm. In other instances, our practice is tracking something worth tracking, but we’re misguided about what it is; so we need to improve our understanding of the phenomena. Sometimes we are clear what we’re tracking, but something else is what we should be or need to be tracking.

Social constructionists are interested in cases where there is a gap between manifest, operative, and target concepts, and in particular, where assumptions about what’s natural are misleading us about what we’re talking about. Constructionists come in many forms, of course, but at least a good number of us argue, concerning certain specific concepts, that contrary to common assumptions, we are tracking something social when we think we’re tracking something natural, and pointing this out is a way of understanding what we’re *really* talking about. So although the constructionist’s analyses may seem revisionary, the proposed revisions in our understanding bring our ideas in better accord with what we have been doing (or should have been doing) all along. This sort of revisionary analysis is surely in keeping with the philosophical goal of talking about what we should be talking about, and being fully aware of what that is.

Given the different projects of analysis and different subject matters for analysis, it is not surprising that philosophers who may appear to be asking the same question are in fact talking past each other. For example, where one philosopher might assume that an adequate analysis must capture our ordinary intuitions, another may take for granted that a priori reflection is likely to be systematically misleading when we are trying to understand the social domain. Recent work on race provides an excellent example of the diversity of

approaches. Some authors are engaged in a conceptual project, attempting to explicate our ordinary understanding of race (Appiah 1996, Zack 1997, Hardimon 2003, Mallon 2004); others are attempting to determine what, if any, natural kind we are referring to by our racial terms (Appiah 1996, Kitcher 1999, Andreason 2000, Zack 2002, Glasgow 2003); others have pursued genealogy (Omi and Winant 1994): and still others are invested in what I call ameliorative projects, raising normative questions about how we *should* understand race, not only how we currently do (Gooding-Williams 1998, Alcoff 2000).

What should we make of these different projects? Should we simply allow that different inquirers are interested in different questions, and nothing can be said to resolve the question what race *really is* or what we mean by “race”? I would not argue that there is one thing that race *really is* or one thing that “we” mean by “race.” Nevertheless, in developing an account of race we should be attentive to our manifest, operative, and target concepts and, if there is a legitimate target notion, have them coincide. It is a mistake, then, for those engaged in conceptual analysis to dismiss inquiries into operative and target concepts, with the thought that only the conceptual project can discover “our” concept. For example, if we discover that we are tracking something that is worthwhile to track in using our racial vocabulary, then even if this is not what we originally “had in mind,” it still may be what we have been and should continue to be talking about.¹¹

There are cases, however, where the different strands of analysis confront each other more directly. Let’s consider again how a gap between manifest and operative concepts arises. Working within an externalist paradigm, the standard case will be one in which the paradigms are projectible onto an objective type, but those whose manifest concept is at issue are typically ignorant or mistaken in some way or other about the type. So, for example, suppose that in a particular community a substantial number of the population take ‘evergreens’ to refer to plants that have needles instead of leaves. In this context, plausibly the manifest concept of ‘evergreen’ will be of plants with needles. However, given externalist considerations about the broader function of the term ‘evergreen’ in that community and plant types, it is reasonable to conclude that the term ‘evergreens’ picks out some broad-leaved plants such as hollies, rhododendrons, some ivies, and the like. Consider another case: suppose (as at MIT) there is a rule that says students may not receive an incomplete for a course unless 80 percent of the work has been submitted. Suppose further, however, that faculty often grant incompletes to students who have submitted less work, but each think that other faculty generally stick to the rule. Here too there will be a gap between what is generally understood by ‘incomplete’ and how the practice distinguishes those who earn incompletes and those who don’t.

In the ‘evergreen’ case, there are compelling reasons to think that those restricting the meaning of ‘evergreens’ to plants with needles are making the

mistake about what evergreens are, given that there are legitimate reasons to distinguish plants that stay green all winter from those that don't, and that members of the community who work with plants and are most familiar with plant types are able to track that distinction with their use of the term. But in the case of incompletes, it is less clear. Once the gap between rule and practice is pointed out, there may be controversy about what an incomplete "really is." The hard-nosed faculty may insist that the rule for incompletes defines what incompletes really are, and any other incompletes were given in error. Let's call this the *strict standards approach*. More accommodating faculty may argue that the actual practice (for example, of giving incompletes to any student who completes some but not all of the work) is what incompletes really are, and the hard-nosed faculty are living in a fantasy if they think the rule is followed. We could call this the *priority of practice approach*. An obvious next move would be to say that the important question is not what incompletes *are*, but what they *should be* that matters, and move the question to an ameliorative inquiry. However, this example highlights that conflict over what we're talking about may turn on how we draw the distinction between paradigms and errors, since of course the hard-nosed and the accommodating faculty don't agree on what the paradigm and mistaken incompletes are.

Defenders of intuition-based (that is, nongenealogical) conceptual analysis are likely to side with the hard-nosed in such cases. Our paradigms, it might be argued, should at the very least conform to our core ideas about how to apply the concept. If we agree on the 80 percent rule for incompletes, then our paradigms for incompletes ought to be those cases in which the rule is followed. If we project the type from the "right" paradigms, our manifest and operative concepts will coincide. The genealogist will insist that our regular practice of granting students incompletes should determine the paradigms, for what we're doing (and have done) with the distinction is what matters. And if we become clear what we're doing, the manifest and operative concepts will coincide.

In many such cases, we face two questions: what policy do we want to promote (or what objective type do we want to track), and what do we want to do with the bit of language we have been using? Do we want to change our policy and keep the same term, change it and introduce a new term, keep the policy and change the term, or keep the policy with the old term? How we proceed is primarily a pragmatic, political, and rhetorical issue. If the term has been long and strongly associated with a particular policy (or type), then it may take substantial work to change what we do with it; other terms are quite malleable in their operations. An important question is whether there ever are cases where the genealogist's (constructivist's) stand that the social category is what we're *really* talking about is the only reasonable option. Although I am inclined to believe that there are such cases, a full defense of this position is not possible here.

Considering the different forms of philosophical analysis, it should now be more clear that in charging that an account of a concept is revisionary, one must do more than show that it violates some ordinary intuitions; moreover, the claim that one's own account captures "our concept," must be explicated and defended by more than trotting out one's own intuitions (or a group of philosophers' intuitions) about how "we" tend to use the concept. Although social constructionist analyses are not what most people "have in mind" when they think about gender or race, it does not follow that they are inadequate even as analyses of *our* concepts, for a genealogical analysis undertakes not to explicate what is in our heads, but rather the constitutive social matrix for the paradigms.

In this discussion, I have done nothing to argue that the best way to account for gender or race in the United States is to undertake genealogy; rather, my aim has been to provide a framework for taking seriously social matrices within the context of philosophical inquiry. I believe that social constructionist accounts of race and gender (and other social categories) are attempts to identify what, among the complex forces and structures of social life, constitute a widespread and enduring source of injustice. Because our manifest concepts of race and gender still tend to be naturalized, it is news, but not conceptual revision, to provide analyses that explain the commonality amongst those of a race and a gender as social. My hope is that greater attention to the gap between manifest and operative concepts will lead philosophers to focus less on our intuitions and more on the role of concepts in structuring our social lives. Philosophical analysis has a potential for unmasking ideology, not simply articulating it.

NOTES

Thanks to Louise Antony, Rachael Briggs, Alex Byrne, Jorge Garcia, Ishani Maitra, Mary Kate McGowan, Ned Hall, Richard Holton, Rae Langton, Marion Smiley, Sarah Song, Ásta Sveinsdóttir, and Steve Yablo for discussing with me the issues raised in this paper. Special thanks to Lawrence Blum, Samantha Brennan, and Anita Superson for comments on an earlier draft. A version of this paper was presented at the Society for Analytic Feminism Conference in London, Ontario, June 6, 2004, and benefited from the discussion.

1. Quine distinguishes different forms of definition, the third being what he calls (drawing on Carnap) "explicative." In giving explicative definitions, "an activity to which philosophers are given, and scientists also in their more philosophical moments . . . the purpose is not merely to paraphrase the definiendum into an outright synonym, but actually to improve upon the definiendum by refining or supplementing its meaning" (Quine 1963: 24–25). "Ameliorative" captures better than "explicative" the sort of project Quine is characterizing as especially philosophical; it is this sort of project

that I've also called "analytical" (Haslanger 2000). Because "analytical" is commonly used to characterize Anglo-American philosophy in general, and because I'm attempting here to introduce a more fine-grained framework, using "ameliorative" rather than "analytical" will sometimes avoid ambiguity. It should be understood, however, that on my view, whether or not an analysis is an improvement on existing meanings depends on the purposes of the inquiry.

2. If there is a single correct definition of *tardy* it should probably generalize over all cases, for example, someone is *tardy* for *X* just in case they arrive after the official starting time for *X* (without an officially recognized excuse), and they are required or expected to arrive on time. But one might argue that even this sort of definition privileges the institutional structure rather than the local practice.

3. It might be worth noting that even the institutional definition given above—which would probably be the one articulated by the staff and students in the school—is not a general definition of *tardy* but only the definition for our school. Even in Cambridge, other public schools have different starting times.

4. Though not always: where institutions are constructed to ensure social justice, the "subversive" meanings are often sites of injustice.

5. I don't mean to suggest here that there is only one manifest concept and only one operative concept. The manifest and operative concepts may vary from context to context. Note also that although it may be tempting to map the manifest/operative distinction onto the more familiar conception/concept distinction, it is not going to do the work needed. For example, in the case of 'tardy' in Isaac's classroom, the manifest concept is the institutional or public one and the operative one is more idiosyncratic. This suggests that in some cases what we think may be more common and public than what we do with language.

6. This matrix is not invisible to the alert. For example, in the 28 March 2004 *New York Times* you can find this headline: "Pollution and the Slippery Meaning of 'Clean.'" According to the article, families in the area of Love Canal "live in neat, new ranch houses and federal officials recently announced that they now consider this notorious symbol of industrial pollution clean. But what does clean mean when the pollutants that rendered Love Canal dangerous to humans remain exactly where they were? In fact, there is no accepted standard, and clean, in practical terms, often means still polluted—but in a different and less dangerous way." Similarly, the term 'clean' when contrasted with 'explicit' in describing rap lyrics doesn't exactly mean what it connotes and serves as a stand-in for a complex social matrix.

7. Although I'm not endorsing the methods of ordinary-language philosophy, the complexity of our use of words in different contexts is something ordinary-language philosophers were well attuned to, and some of their methods and ideas are tremendously valuable for genealogy.

8. Because the terminology of 'natural kind' is used in several different ways, it is helpful to make a few distinctions. The term 'kind' is sometimes used to classify substances, in the ordinary case, (physical) objects. Substances can be classified according to their essence; kinds consist of groups of objects with a common essence. For example, tigers constitute a kind of thing because each tiger has essentially a certain cluster of properties that define the kind. On other occasions, the term 'kind' is used to refer to

what are sometimes called types. A type is a group of things, sometimes substances, but possibly nonsubstances, that has a certain unity. This unity need not be a matter of sharing essential properties: red things constitute a type (their unity consists in their all being red), even though redness is seldom an essential property of the things that have it. Unity seems to come in different degrees. The things on my desk might be thought to constitute a weak sort of type (they have in common the fact that they are on my desk), and at the limit there are highly gerrymandered sets of things that don't have any unity at all and so don't constitute a type.

One way to think about the unity of types is in terms of similarity between the members. We can distinguish different sorts of types by distinguishing axes of similarity. Exactly six foot tall human beings are a natural type because the commonality between the members is natural (species and height); high school graduates are a social type because the commonality between the members is social. Both of these types are (metaphysically) objective, however, in the sense that the commonality between the members lies in properties of the objects (or relations between them), and not in their relationship to the speaker or cognizer. How to draw the line between social and natural types is difficult and not one I address here. I rely on background understandings and familiar cases. However, it is important to keep in mind that as I am using the terms, the distinction between objective and nonobjective kinds or types is importantly different from the distinction between natural and social kinds or types.

9. The third sort is supported by the idea that inquiry can seem to converge on an idea or concept that we seemed to have in mind all along, even though no one, even the best minds, could have explicated it. So, for example, in the history of logic, math, and science, it is plausible to maintain that certain experts were “grasping a definite sense, whilst also failing to grasp it ‘sharply’” (Peacocke 1998, 50). I assume that these can fall within objective type externalism because at least there are some paradigms that fix the reference in question; they don't seem to qualify for social externalism because there isn't yet a standard linguistic usage.

10. I put “our” and “we” in scare quotes to indicate that there may be significant contextual variation, or at least there will be room for contestation.

11. It might be useful to see this by analogy with other terminological developments in science. Although our understanding of and even our definition of ‘atom’ has changed over time, it is plausible that there is something worthwhile we have been and continue to be talking about.

REFERENCES

- Alcoff, Linda. 2000. Is Latina/o identity a racial identity? In *Hispanics/Latinos in the United States*, ed. J. J. E. Gracia and P. DeGrieff. New York: Routledge.
- Andreasen, R. O. 2000. Race: Biological reality or social construct? *Philosophy of Social Science* 67 (Proceedings): S653–S666.
- Appiah, Kwame Anthony. 1996. Race, culture, identity: Misunderstood connections. In *Color conscious: The political morality of race*, ed. Kwame Anthony Appiah and Amy Gutmann. Princeton, N.J.: Princeton University Press.

- Burge, Tyler. 1979. Individualism and the mental. *Midwest Studies in Philosophy* 4: 73–121.
- Fraser, Nancy. 1989. Foucault on modern power: Empirical insights and normative confusions. In *Unruly practices: Power discourse and gender in contemporary social theory*. Minneapolis: University of Minnesota Press.
- Gooding-Williams, Robert. 1988. Race, multiculturalism and democracy. *Constellations* 5 (1): 18–41.
- Hacking, Ian. 1999. *The social construction of what?* Cambridge, Mass.: Harvard University Press.
- Hardimon, Michael. 2003. The ordinary concept of race. *Journal of Philosophy* 100 (9): 437–55.
- Haslanger, Sally. 1995. Ontology and social construction. *Philosophical Topics* 23 (2): 95–125.
- . 2000. Gender and race: (What) are they? (What) do we want them to be? *Noûs* 34 (1): 31–55.
- Kitcher, Philip. 1999. Race, ethnicity, biology, culture. In *Racism*, ed. L. Harris. Amherst, N.Y.: Humanity Books.
- Kripke, Saul. 1980. *Naming and necessity*. Cambridge, Mass.: Harvard University Press.
- MacKinnon, Catharine. 1987. *Feminism unmodified*. Cambridge, Mass.: Harvard University Press.
- Mallon, Ron. 2004. Passing, traveling and reality: Social construction and the metaphysics of race. *Noûs* 38 (4): 644–673.
- Mills, Charles. 1998. *Blackness visible: Essays on philosophy and race*. Ithaca, N.Y.: Cornell University Press.
- Omi, Michael, and Howard Winant. 1994. *Racial formation in the United States*. New York: Routledge.
- Peacocke, Christopher. 1998. Implicit conceptions, understanding, and rationality. *Philosophical Issues* 9: 43–88.
- Putnam, Hilary. 1973. Meaning and reference. *Journal of Philosophy* 70 (19): 699–711.
- . 1975. The meaning of ‘meaning.’ In *Mind, language and reality*. *Philosophical Papers*, Vol. 2. Cambridge: Cambridge University Press.
- Quine, W. V. O. 1963. Two dogmas of empiricism. In *From a logical point of view*. New York: Harper and Row.
- Zack, Naomi. 1997. Race and philosophic meaning. In *Race/sex: Their sameness, difference, and interplay*, ed. Naomi Zack. New York: Routledge.
- . 2002. *Philosophy of science and race*. New York: Routledge.