

1 The Race Debate

This book is an attempt to grapple with a problem: the concept of race seems irredeemably corrupted but in some ways too valuable to do without. Now of course this isn't the only race-related problem worthy of our attention. There are the venerable and important questions of racism, and of affirmative action and reparation. And questions of identity, and of the phenomenology and existential significance of race, have been resurgent topics for a couple of decades now. But the questions that motivate this volume excite different curiosities.

In 1897, W.E.B. Du Bois, faced with the question of whether people of African descent should assimilate or carve out a distinct community in the United States, and indeed on the world stage, gave his seminal lecture, "The Conservation of Races." He argued, with characteristic power, not only that this population constituted a race, but also that it had something of a unique mission in the history of humankind, and thus he concluded that the elimination of racial differentiation would be a grave mistake. In the decades that followed, the soundness of racial thinking mostly became a topic to be studied by social and natural scientists, not philosophers. Indeed, with a few notable exceptions, race-thinking was a largely dormant topic in philosophy until the 1980s. At that point, in a now-classic article, Kwame Anthony Appiah (1985; 1992) reexamined Du Bois' conservationism. Dispatching Du Bois' claims like so many badly outdated fashions, Appiah began a series of arguments in defense of the position that race is an illusion unworthy of our credence.

The desire to leave race behind is, of course, a dominant theme of the modern United States. In its least contestable form, it is the sentiment expressed in Martin Luther King, Jr.'s hope that his children be judged not "by the color of their skin but by the content of their character." But *racial eliminativism* makes a stronger claim than that. According to one political version of eliminativism, we should eliminate racial categories from all or most state policies, proceedings, documents, and institutions. Californians rejected such a proposal when in 2003 they voted to defeat the "Racial Privacy Initiative" (Proposition 54), which

2 *The Race Debate*

would have prevented most government agencies from collecting most types of racial data. Not one to be found in lock-step with Californians, George Will (2003), the prominent conservative columnist, has called for removing racial categories from the census. Sometimes, as with Will, political eliminativism is motivated not only by the claim that the way we think about race might be incoherent, but also by the rationale that eliminating racial categories will undermine other policies, such as affirmative action, which presuppose race. Indeed, in case there was any question, the brouhaha over eliminativism was that same year declared “a national debate” by the front page of the *New York Times* (Nov. 9).

A second, more sweeping form of eliminativism is the public version. Public eliminativism advises that we get rid of race-thinking not only in the political sphere, but in the entirety of our public lives, so that we neither assert nor recognize one another’s races. Finally, there is global racial eliminativism. The goal of this view is for us to eventually get rid of race-thinking not only in the political or even public world, but altogether. That is, even in our most private inner moments, race-thinking should go the way of belief in witchcraft and phlogiston: a perhaps understandable but hopelessly flawed, antiquated way of making sense of our world, a way of making sense that has no place in our most sophisticated story about The Way Things Are.

Now, in the wake of eliminativism’s rise, several respondents have tried to update and defend Du Bois’ basic position that race-thinking is worthy of conservation. These *conservationists* argue that, for various reasons to be examined within these pages, eliminating race-thinking would be a serious error. Thus, to take them out of order, the first of four main questions to be asked here, the question that will set much of our agenda in crucial respects discussed below, is

The Normative Question: Should we eliminate or conserve racial discourse and thought, as well as practices that rely on racial categories?

Those, anyway, are the conventional options. But, as is often the case with convention, this set of choices unnecessarily presents us with too few options. Or so I will argue. To turn over my first card, the normative position I will advocate is neither that we should out-and-out eliminate race-thinking, nor that we should wholeheartedly conserve it, but that we should *replace* racial discourse with a nearby discourse. The basic idea to this position—what I will label *racial reconstructionism*—will be that we should stop using terms like ‘race,’ ‘black,’ ‘white,’ and so on to purport to refer to biological categories—

as we currently use them. Instead we should use them to refer to wholly social categories.¹

It is worth pausing for a moment to emphasize who ‘we’ are here. We are neither philosophers in particular nor academics in general. The Normative Question is whether *all of us*—everyone in our linguistic community—should keep or abandon racial discourse. Racial reconstructionism says that all of us should reconstruct our racial discourse. (Arguably racial discourse operates differently in different communities, so I will focus particularly on my linguistic community, which comprises competent English speakers in the United States. That said, I suspect that the arguments found below are relevant in many other communities as well.)

Whether we should be eliminativists, conservationists, or reconstructionists depends on two main considerations. First, a clutch of particularly salient evaluative considerations bear on this question: is racial discourse morally, politically, or prudentially valuable? For instance, if someone wants to be identified in a certain way, we arguably have a moral obligation—one that in some contexts can be overridden, to be sure—to identify them in this way. Obviously, racial identities are key components of some people’s self-conceptions, so moral value will have to be addressed here. A political question relevant to our discussion is whether race-thinking enables important policies for redressing racial injustices, or whether, as the biologist Joseph L. Graves (2001, 11) maintains, “the survival of the United States as a democracy depends on the dismantling of the race concept.” Less bold, but equally pressing and more common, is the political *and* moral claim that getting rid of race-thinking is part of a program of getting rid of racism (Appiah 1996, 32; Graves 2001, 200). Finally, abandoning race-thinking might be prudentially bad because doing so would disintegrate one’s individual identity; or it might be prudentially good because it allows us to pursue relationships that are difficult to pursue in a race-conscious world. And, of course, sometimes all of these values are thrown together into one mess. For instance, Graves (2001, 199) proposes an item that potentially impacts the putative political, prudential, and moral value of eliminating race-thinking: doing so will foster economic growth.

Before we get to those kinds of concerns, though, note a second relevant issue. If race is not real, then that generates one reason to get rid of racial discourse; if it is real, then that provides at least one reason to retain it. This presupposes a principle of epistemic value: if our beliefs should be sensitive to available evidence, then it is bad both to believe

1 Following convention, I use single quotation marks when mentioning words and phrases, and small caps to name concepts.

4 *The Race Debate*

in something that evidently doesn't exist, and to pretend that something that evidently does exist doesn't. Other things being equal, you shouldn't believe that an invisible goblin is typing these words for me so that I can relax and enjoy a beer. And other things being equal, you shouldn't pretend that the moon doesn't exist. If you're with me on this—if you agree that, other things being equal, we shouldn't believe in things that evidently aren't real and that we should believe in things that evidently are real—then you're with me in attributing importance to the second main issue to be discussed in this book, namely

The Ontological Question: Is race real?

When I first mention to civilian friends and students that many academics think that race is nothing but an apparition, one common reaction is incredulity. To such a way of thinking, the fact that each of us has a race, or multiple or mixed races, is unassailable. Any departure from conventional wisdom here might make academics appear to be unglued from the real world by sheer force of theoretical peculiarity. Whether or not the glue still holds will be an overarching theme in this book, as one of my main concerns—a concern that, I will argue, has been problematically ignored by many (myself included, at times)—is to account for, or at the very least confront in a richly informed way, commonsense thinking about race.

Though we will see below that commonsense thinking about race is in fact strikingly complex, one fairly predominant element of the folk theory of race is that races are biological entities. Now there is more than one way that race might be biologically real. According to one understandable line of thought, we have skin colors and hair textures and facial features—we have, as biologists like to say, *phenotypes* (roughly, the macro-level expressions of our *genotypes*, our genetic makeup). If we can classify these phenotypes in a biologically kosher manner, then this is one way in which race might be biologically real. As it features something that can be superficially read off of the way we look, I'll call this view the *superficial theory*. Another way in which race might be biologically real is not in terms of what we look like, but in terms of the genetic material that significantly determines what we look like—a theory we can call *genetic racial realism*. And then there is the source of our genetic material, namely our ancestry. So a currently popular wave of biological racial realism—*populationism*—holds that races are breeding populations or clusters of breeding populations, populations whose intra-group reproductive rate is sufficiently higher than their rate of reproduction with other populations, thereby ensuring genetic distance (and, usually, phenotypic difference) from each other over multiple generations.

As we will see, all three biological accounts of race have problems; but there is another branch of racial realism. Many contemporary realists, taking inspiration from Du Bois, maintain that race is not ultimately about biology at all. Instead of being a biological kind of thing, race is, on this alternative theory, socially constructed but real nonetheless. That is, race is real as a *social* kind of thing. This view, which I'll call *constructivism*, holds that just as journalists or doctors are real but socially constructed kinds of people, so racial kinds of people are real but socially constructed—racial groups are real groups that have been created by our social practices, rather than by some biological process. Thus there are several different types of realism one might adopt (see Figure 1.1).²

Anti-realists generally think that race is not real because race purports, but fails, to be a biological kind. (Strictly speaking, though, this specific route to anti-realism is not required to be an anti-realist.) Obviously, each kind of realism is inconsistent with this anti-realist thought in its own way. Biological racial realists argue that anti-realism is wrong, because there is (they say) a biological reality to race. Alternatively, constructivists argue that race doesn't need to be a biological kind of thing to be real; instead, it's a socially constructed kind of thing. Thus on my way of defining the various theoretical positions, 'constructionism' names the view that the *idea* of race has been socially

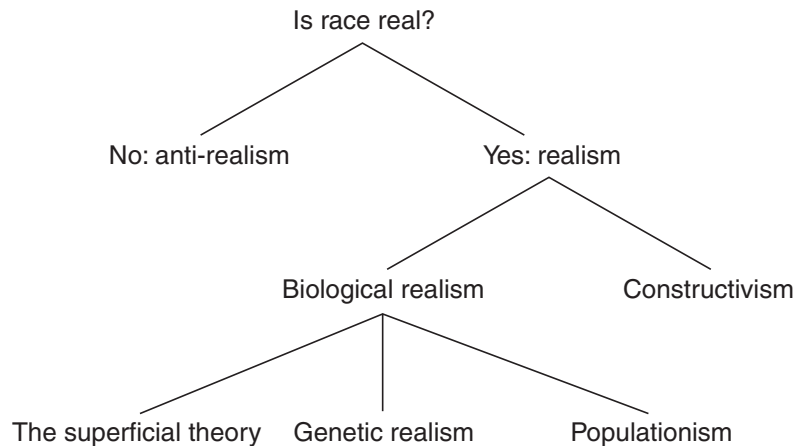


Figure 1.1 The landscape.

2 What I am here calling 'racial realism' others sometimes call 'racialism.' Racialism, however, is often taken to be a more robust view of race (such that, say, phenotypic features are correlated with intellectual capacities). 'Racial realism,' which just holds that there really are races, is meant to be neutral between these more robust and other, more austere, accounts of race. I should also note that, in this way, I am also using 'real' in what is sometimes called its *minimalist* sense, according to which something is real just in case it exists. The use of the minimalist sense, which creates conceptual room for constructed or response-dependent entities, is required if we are to take constructivism seriously.

6 *The Race Debate*

constructed, and to that extent it is neutral between anti-realism and the view that there are socially constructed races. This latter view is what I call ‘constructivism,’ and given the theoretical space it occupies, in order to know whether race is real we now have to answer the more basic question of what race is *supposed* to be. Is it supposed to be a biological kind (and if so, of what sort), or a social kind? Put somewhat differently, what are we purporting to talk about when we use words like ‘race’? This is the third main issue of contention that will be examined here:

The Conceptual Question: What is the ordinary meaning of ‘race,’ and what is the folk theory of race?

Rather than asking whether there is something in the world that matches up with our race-talk, this question just targets our race-talk itself: what kinds of things are we purporting to talk about when we talk about race? Are we trying to talk about real scientific kinds, as we do when we talk about, say, water or gold? Or are we purporting to talk about illusory kinds, as we did with, say, witches or phlogiston? Or, finally, could we be talking about some element of the socially constructed world, as when we talk about money or journalists or universities, which have no place, in and of themselves, in the world studied by natural scientists? Although, as we will see, what is in the world can sometimes help determine the meanings of our terms, the Conceptual Question is in the first instance a question about our racial discourse: What kinds of things are we *purporting* to refer to when we talk about race? At its core, this is simply a question about the *meaning* of ‘race’ and cognate terms. So the core part of the Conceptual Question is semantic. At its periphery, this question asks not about the ordinary meaning of ‘race,’ but about the folk *theory* of race; so the other part is folk-theoretical.³

It’s hard to overstate the importance of this question. If racial discourse does not purport to refer to a biological kind, then it will be a non-starter to argue that races are not biologically real. For, assuming that some things, such as universities or newspapers, are real not as biological things but instead as social things, race might be non-biologically real, too. If, however, racial discourse *does* purport to pick out biological categories, then when constructivists tell us that race is a

3 I should note, then, that I am using ‘conceptual’ in an exceptionally broad way. Concepts are just the meanings of terms, but as used here, ‘conceptual’ covers not only concepts, but also folk theory and belief, or conceptions. This distinction doesn’t matter now; it will matter below, though, starting in Chapter 2.

social kind, they will be the ones who are talking about something else besides race, and their position would be the one that is irrelevant. It would be comparable to a debate about whether there were any real witches in colonial Salem, in which I insisted that there were, because some of Salem's residents practiced Wicca. Aside from it being factually false that they practiced Wicca, you'd legitimately have a more basic conceptual complaint: those 'witches' are not the witches you're talking about when you deny that there were any witches in Salem. In that context, by 'witch' one means the kind of person who casts spells and cavorts with the devil, so it is irrelevant whether anybody in Salem was practicing Wicca.⁴ In a parallel kind of way, once we know what race is *supposed* to be, we can figure out whether there is, in fact, any such thing.

That's all by way of saying that the Conceptual Question is dialectically important: if we want to figure out an answer to the Normative Question, it seems as though we're going to have to try to answer the Ontological Question, which means having to answer the Conceptual Question. Without minimizing this dialectical importance, we also should not forget that the conceptual truth about race has a substantial impact on the real world. Lucius Outlaw makes the point powerfully in the course of examining the nature and function of our racial categories:

this is more than an issue of philosophical semantics in racially hierarchic societies which continue to engage in efforts to promote and sustain forms of racial supremacy. In this context, racial categories take on the various valorizations of the hierarchy and affect the formation and appropriation of identities as well as affect, in significant ways, a person's life-chances.

(Outlaw 1996a, 33)

And lurking behind the crucial Conceptual Question is one final issue. The Conceptual Question asks what racial discourse purports to talk about: anti-realists, such as Appiah (1996) and Naomi Zack (1993, 1995, 1997, 2002, 2007), think that ordinary racial terms (erroneously) purport to refer to some sort of interesting biological reality, while

4 Here I depart from Appiah (2007, 38–39), who holds that, while the identity *witch* has different criteria of ascription in the two contexts, there is a live question as to whether we should “give up the concept” or preserve the *term* ‘witch’ for picking out practitioners of Wicca. I think this way of framing things obscures a more natural reading of the linguistic and conceptual terrain: we can simply say that the term ‘witch’ is ambiguous, such that on one meaning it purports to refer to a supernaturally gifted friend of the devil and on the other to practitioners of Wicca. Indeed, we should say this *because* of the two contexts’ radically different criteria of ascription. In this way, you don’t have to make a hard call about giving up *the* concept; you simply give up one concept and keep the other, while the term remains the same.

many of their opponents think that they purport to refer to some sort of social reality. But, then, if we're going to try to figure out what racial terms purport to refer to, we need to know *how* to figure that out. That is, we must also answer

The Methodological Question: How should we identify the folk concept and theory of race?

As we shall see, one answer to this question is that in order to identify the folk concept of race, we should look at how experts have historically used racial terms. To turn my second card face up, I will argue that this methodology is, by and large, misguided. Instead, I will maintain that, for the most part, we should focus our attention squarely on how racial terms are used in contemporary mainstream discourse. Some people agree with that approach, and then proceed to engage in personal reflection—they reflect from the armchair, as we say—about the nature of contemporary folk racial discourse. I will also argue that the armchair-based approach is, while useful to an extent, insufficient. As an alternative, I adopt what I call the ‘experimental approach,’ which holds not only that the meanings of racial terms are, for our purposes, at least partially fixed by commonsense, but also that we should inform our analysis of folk racial discourse with data gathered from actual empirical research conducted in a manner consistent with the practices of the social sciences. To be sure, I, like many, accept that we can also identify some of the content of our racial concepts while comfortably ensconced in the armchair. But even the data gathered from such armchair expeditions must be consistent with the empirical data. So that’s how I answer the Methodological Question.

Let me reveal the rest of my hand at this point. After setting out, in Chapter 2, some conceptual limits from the armchair and defending my methodological approach in Chapter 3, I turn in Chapter 4 to completing my answer to the Conceptual Question by looking at recent provocative empirical research. Once that conceptual groundwork is laid, we will be in a position to address the Ontological Question of whether race, as defined in that relevant sense, is real. In Chapters 5 and 6, I will argue that it is not. Thus Chapters 2 to 6 constitute an extended argument for the claim that race is not real. But when we consider the Normative Question in light of that claim, we’re left with something of a puzzle; for, as I suggest in Chapter 7, we’d be poorly advised to simply get rid of racial discourse. My attempt at a solution is, again, to argue that instead of merely conserving or eliminating racial discourse, we need to replace it with a nearby discourse. Now there are various nearby discourses that are candidates for replacement.

		Racial ontology	
		Anti-realism	Realism
Racial politics	Eliminativism		
	Anti-eliminativism	Reconstructionism (substitutionism)	Conservationism

Figure 1.2 The theoretical options.

To mark the general theoretical genus of which this reconstructionism is a species, I generally call replacement theories *substitutionism*. So reconstructionism is one particular kind of substitutionism (see Figure 1.2).

Those, then, are the issues engaged in this volume. In brief, the question of whether we should eliminate or conserve race is tangled up in the thorny question of whether race is real or illusory. Whether or not it is real, in turn, depends on what race is supposed to be—on how we use racial discourse—which requires us to do some analysis of ordinary racial concepts; this analysis, in turn, requires us to answer a methodological question about how to analyze our folk concepts and theories. Taken together, I am calling the ongoing disputes over how to best answer these four questions *the race debate*. Should my answers carry the day, the ultimate payoff is the normative proposal of racial reconstructionism. At this point, then, we can preview the three arguments that jointly comprise the overall master argument of this book. Philosophers have found that readers tend to find books that present arguments in lists of displayed premises and conclusions ‘page-turners,’ so without further ado:

The Conceptual Argument (Chapters 2–4)

- 1 The discourse relevant for the race debate is folk racial discourse, and analysis of folk racial concepts (and conceptions) should be informed by empirical study.

- 2 According to the most plausible, empirically informed analysis, folk racial concepts (and conceptions) are composed of biological elements and sometimes also social elements.

Thus,

- 3 The relevant racial concepts (and conceptions) are composed of biological, and sometimes social, elements.

The Ontological Argument (Chapters 5–6)

- 3 The relevant racial concepts (and conceptions) are composed of biological, and sometimes social, elements.
- 4 There are no biological races; and purely social ‘races’ aren’t races in the relevant sense.

Thus,

- 5 Race, in the relevant sense, is not real.

The Normative Argument (Chapter 7)

- 6 What we should do with our racial discourse is ultimately a matter of which of the various proposals—eliminativism, conservatism, and substitutionism—best satisfies various normative (moral, political, prudential, and epistemic) constraints.
- 7 Given that (5) race is not real, adopting racial reconstructionism is the best way to satisfy all of the normative constraints.

Thus,

- 8 We should adopt the policy of racial reconstructionism.

That’s the set-up. But before moving on to defend these arguments, I want to address one last preliminary issue. Some—the *exclusionists*, as I will call them—think that the only business philosophers have in the race debate is to tackle normative or evaluative problems: we do ethics and politics, but we should leave the biology to the biologists, the anthropology to the anthropologists, and so on. Now I of course agree that I have no business doing biology, but I think it’s too quick to say that the race debate doesn’t involve *non-normative* philosophical

questions, in particular ontological and conceptual questions. I favor the non-exclusionary thesis that addressing the ontological and conceptual questions is both a philosophical job and a job worth doing. So I want to finish this chapter by driving home why, and in what respects, the questions taken up in this book are both philosophical and well motivated. I don't hope to *answer* the conceptual or ontological questions just yet. Rather, the current goal is to ascend to the meta and determine, on the assumption that we want to answer the Normative Question, which other questions must be answered, and in particular which must be answered by doing some philosophy.

The exclusionists are surely right to insist that the Normative Question is of independent importance and that any answer to the Ontological Question does not by itself settle the issue of what policy we should adopt.⁵ For even if race is real, racial discourse could be harmful; and even if race is an illusion, racial discourse might serve some important interests. But the exclusionists go beyond merely advocating that we give special attention to distinctively normative concerns. Ron Mallon (2006, 551) adds that the Normative Question is *the* important question, and that the ontological and conceptual questions should be "left behind." Anna Stubblefield (2005, 73) similarly considers attention given to the Ontological Question not only "unhelpful," but actually also a hindrance to answering the Normative Question.⁶ And, whether or not the Ontological Question is helpfully relevant to the Normative Question, Bernard Boxill (2004) maintains that it properly falls under biologists' area of expertise and so is not a philosophical question at all.

To begin to see why one might go exclusionist, briefly recall the dialectic discussed above: it is sometimes held that (in part) because race is not real, we should eliminate racial discourse; and race is not real, some say, because races are supposed to be certain biologically grounded collections of people, which as it happens turn out not to exist. Their opponents, the conservationists, often defend keeping racial discourse partly on the grounds that race is real. So the broader dialectic is such that the normative dispute over whether to eliminate racial discourse often—though not always—boils down in part to an

5 Non-exclusionists often seem to tacitly agree with this; for some who have explicitly agreed, see Glasgow (2006); Kitcher (1999, 90; 2007); Sundstrom (2002a); Taylor (2000, 2004).

6 What Mallon and I call 'the normative question' Stubblefield in many places calls 'the moral question.' 'Normative' seems like a more apt term, since diverse realms of normativity besides just the moral bear on the issue, including the political, the prudential, and the epistemic. Or perhaps Stubblefield intends us to focus, even more exclusively, on *just* moral issues, and to not attend to political, prudential, and epistemic issues at all. My reasons against making our focus so very narrow are presented below.

ontological dispute: if race is an illusion, we should abandon racial discourse, while if race is real, we may keep it (barring overriding evaluative concerns). And, again, the ontological debate itself partly reduces to a further dispute over the meaning of ordinary racial terms. Thus the broad orientation of the race debate is that the normative is held to depend on the ontological, which in turn is held to depend on the conceptual. Mallon calls this dialectical move, of defending a normative position partly on the basis of ontological and at bottom semantic (and, I would add, folk-theoretical) theses, “the semantic strategy.”⁷

Part of Mallon’s resistance to this strategy stems from the fact that amidst the din of disagreement there is an important and robust set of claims that everyone agrees on, which he rightly calls the “Ontological Consensus.” It includes such propositions as that “[t]here are no biobehavioral racial essences,” that people use such factors as skin color and ancestry as criteria for racial identification, and that racial classification has had oppressive effects (Mallon 2006, 545). The list is longer, but what is remarkable about it is how widespread the agreement is. I certainly don’t wish to dispute that point: we should agree with Mallon that this common ground should not be neglected. Nevertheless, I do want to defend the value of tackling the ontological and conceptual questions against the exclusionist arguments to the contrary. So let us take those arguments one at a time.

One exclusionist argument is based on a sensitivity to disciplinary boundaries and a healthy respect for expertise. Here is how Boxill makes his case:

philosophers are probably not in the best position to prove that there are no races. Full-time biologists seem to be in a better position, given that by ‘race’ we mean here *biological* race, namely a group of individuals defined biologically, like a breed or a subspecies.

Appended to this claim is a footnote: “Recently, ‘race’ has sometimes come to have a different meaning, as referring to a social construct. The existence of race as a social construct is not controversial” (Boxill 2004, 209).

⁷ Mallon and others, including my past self (Glasgow 2006), treat it as *just* a semantic dispute, but sometimes the contentious matter is not about the *meaning* of folk racial terms, but about the *folk theory* of what race is. To wit, two parties can use a shared concept, such as RACE, but disagree about the substantive nature of race, e.g., as to whether it is a natural or social kind. Hardimon (2003) has shed light on this distinction by noting that what is often claimed to be the *concept* of race is actually an account of the ordinary *conception* of race. Again, this distinction will receive fuller treatment in Chapter 2.

Boxill is surely correct that if the reality of race were just about the biological facts, then (most) philosophers should step aside. Still, it would overstate the implications of this claim to infer from it that philosophers aren't the ones to argue that there are no races. For while biologists certainly give us the biological facts, there are reasons—at least four reasons—why the answer to the question of whether race is real is underdetermined by facts supplied by other disciplines, including biology.

First, we need to justify what Boxill takes as given, namely that “by ‘race’ we mean *biological* race.” (Since he takes this as given, perhaps Boxill could agree on this.) According to anti-realists like Appiah and Zack, we do mean this, but according to constructivists we do not. So there is an important disagreement over what we mean by ‘race’ that has to be sorted out before we can hand the discussion over to biologists, and, of course, conceptual analysis is a philosophical task.

This point is not just idle disciplinary defensiveness; it allows us to identify errors that show a strong tendency to afflict all parties to the race debate. For example, as part of his attack on the view that races are real, Graves (2001, 5) writes that “[t]he term ‘race’ implies the existence of some nontrivial underlying hereditary features shared by a group of people and not present in other groups.” While Boxill is right that most philosophers have no business questioning the science behind Graves’ attack, it is our business to question the definition of race that Graves presumes. He gives no argument for his definition, although he does reassert it several times, such as when, considering *The American Heritage Dictionary*’s six-part definition of race, Graves insists that only the fifth part, on heredity, is “a correct scientific definition of race” (Graves 2001, 6) (which ends up having no real-world referent, he argues). But this approach assumes what it needs to prove, namely that ‘race’ is defined in the scientific way he thinks it is, for, if those who see races as social kinds are right, then Graves’ semantic presumption in favor of a biological analysis is wrong. And if that analysis is wrong, then, no matter how compelling his argument that there is no biological reality to race, this won’t establish that race isn’t real, because race is (if the constructivists are right) some *other* kind of thing, some *social* kind of thing. In this way, the semantics of race can, at least in principle, render biology irrelevant.

In contrast to Graves, Vincent Sarich and Frank Miele (2004, 14–15) defend racial realism, but in so doing they make a mistake that runs exactly parallel to the one committed by Graves: they also assume that the dictionary definition of ‘race’ (which they take from the OED: “a group of persons connected by common descent’ or ‘a tribe, nation, or people, regarded as of common stock’”) reflects the “commonsense” definition, and that this commonsense definition is self-evident. But simple fiat cannot be a substitute for thorough analysis. So whether one

makes Graves' faulty presumption in favor of racial anti-realism or Sarich and Miele's faulty presumption in favor of realism, either way there is a faulty presumption, namely that one can simply assert the ordinary definition of 'race.' Indeed, if it were as obvious as they suppose, then Graves and Sarich and Miele should converge on the same definition. Tellingly, they don't.

The more general point is that any argument that utilizes biological facts for or against the reality of race requires that 'race' be defined in a way that makes those facts relevant. And so long as we are trying to characterize the *folk* concept of race, rather than simply stipulating a definition that as a consequence of being stipulated may not engage the race debate at all, any such definition itself needs defending. Giving, and arguing for, those definitions is where philosophical work is animated. That is to say, to analyze the ordinary concept of race, which may or may not pick out something real, we need to bring philosophical tools to bear—counterexamples, thought experiments, fine-tuned analysis, and so on. These tools will then shape the analysis of both the core concept of race and the broader folk theory of race that realists will seek to vindicate and anti-realists will seek to eviscerate. Note: to say that it is a philosophical task is not to say that the best analysis will not be informed by empirical psychological data regarding how we think about race. Indeed, if Chapter 3 is correct, the reverse is true. Furthermore, it of course need not be *philosophers* who do the work of conceptual analysis (although philosophers presumably have the kind of training which facilitates that work, just as biologists have training that facilitates doing biology); but it is *philosophical* work, no matter who does it.

A second uniquely philosophical issue is whether breeds or subspecies are real. Even if biologists could come up with a sensible division of the human species that includes human races as subspecies, it is an open philosophical question whether subspecies count as real. So in this way, even when the biological facts are centrally relevant, the door is not altogether shut on philosophy. Third, it is not uncontroversial that race exists as a social construct, contra Boxill. While it might be uncontroversial that racial *discourse* exists as a social construct, the claim that *race* is a real social construct is the kind of premise that folks like Appiah, Zack, and myself would reject on semantic grounds: for us, whatever might exist as a pure social construct is not, we argue, what ordinary people call 'race.' So those who want to defend the social reality of race must show how it is semantically kosher to fold race into a wholly social reality—how, that is, they are not talking about something else other than race when they talk about social constructs. The general point here is parallel to the general point made about biological realism and anti-realism: if social facts are going to be marshaled in support of the reality of race (as

ordinarily talked about, that is, as the potential target of elimination or conservation in *public* racial discourse), ‘race’ will need to be defined in such a way that social facts are not irrelevant, such as would be the case if RACE turns out to be a purely biological concept. Fourth and finally, some might deny that non-scientific, social kinds are real at all.

I hasten to emphasize that I do not mean to suggest that any of these questions cannot be decisively answered. All I mean to argue currently is that they are not closed and that obtaining answers to them will come from doing philosophy. The second and fourth of these areas of dispute fall in the domain of the philosophy of biology and metaphysics, while the first and third open up conceptual questions about race. So it appears that some important questions pertaining to the reality of race are, in fact, philosophical. With that, we can now turn to a different kind of exclusionist argument, which holds not that questions about the reality of race fall outside of philosophy’s domain, but that focusing on the reality of race in relation to the Normative Question is a bad idea. Instead, the critics say, the question of conserving or eliminating ordinary racial discourse and race-based practices should be decided independently of the ontological and conceptual questions.

We would have one reason to avoid pursuing the ontological and conceptual questions if the semantic strategy were “obfuscating,” and, according to Mallon (2006, 548), it is obfuscating, because “it makes a philosophical debate over the reference of racial terms and concepts appear as a genuine metaphysical disagreement about what is in the world.” However, the semantic strategy need not be seen as concealing a semantic dispute within a superficial, merely apparent ontological one. Rather, it may be understood such that it presents a *genuine* ontological issue as coexisting alongside of, and in significant part because of, a semantic issue. Indeed, far from obfuscating, some anti-realists are transparent that they aim to answer the Ontological Question by showing that the world doesn’t match up with our discourse (e.g., Appiah 1996; Glasgow 2006; Zack 2002). Furthermore, this kind of argumentative gambit is nearly ubiquitous in philosophy. Debates over the existence of free will, for instance, sometimes trace to deeper disputes concerning the meaning of ‘free will.’ Or, for a further example, Walter Sinnott-Armstrong (2006, 7) observes that in meta-ethics, “[o]ntological conclusions are often drawn from semantic premises.” In this way, just because an ontological issue partly reduces to a semantic issue, we should not conclude that there *is no* ontological issue.

A second objection to the semantic strategy is not that it obfuscates, but rather that with respect to the independently important Normative Question—*Should* we conserve or eliminate racial discourse?—the correct metaphysics and semantics of race might just be *beside the point*. Now it is certainly true that the Normative Question does not

fully reduce to the Ontological Question. This creates theoretical space for people like me to endorse anti-realism but avoid eliminativism and for others to hold that race-thinking is dangerous, even if race is real (Boxill 2004). On these grounds, Boxill (2004, 224) concludes that what is important to the Normative Question seems to be whether racial discourse is valuable, not whether race is real. In a similarly exclusionist vein, Stubblefield (2005, 80) thinks that we should challenge “the assumption that the morality of taking race into account is dependent upon whether or not race is somehow real.”

Now as Mallon (2006, 549–550) recognizes, when asking whether race-talk is valuable, we have to consider not only its moral, prudential, and political value, but also its epistemic value. And if we can agree that it is epistemically bad to believe in something that evidently doesn’t exist, and if our ordinary race-talk encourages us to believe in something that evidently doesn’t exist, then the epistemic consequence of anti-realism is that we should get rid of race-talk. But, of course, whether our racial discourse concerns something that doesn’t exist is the Ontological Question, which depends, as we have seen, on the Conceptual Question. Similarly, if on the correct semantics (whatever they may be) races do exist, then, if we can also agree that we should not pretend that what is real is *not* real, we have that much epistemic reason to keep race-talk around. In this way, part of the project of answering the Normative Question is determining the epistemic value of racial discourse, and determining the epistemic value of racial discourse depends on determining whether race exists (the Ontological Question), which depends on determining what race is supposed to be (the Conceptual Question).

Of course, epistemic value might ultimately be outweighed by some moral, political, or prudential value that is assessable independently of the reality of race. But, first, I will argue that we have the ability to avoid trading away *any* of these values, and, second, even if we needed to determine such a weighting, we’d not only have some fairly complicated moral, prudential, and political issues to sort out, we’d also have to identify the relevant epistemic harms and benefits. That is, rather than simply sidelining some of the relevant values, we need to have all of them before us. So while Stubblefield (2005, 11) holds that we “end up going around and around on the question of what race is and whether it is real and never get to the heart of the matter, which is the moral question,” we should instead recognize that the morality is not the only relevant axis of value.

But even if the Conceptual and Ontological questions are neither obfuscating nor beside the point, Mallon (2006, 548) presents a third objection: the semantic strategy is “ineffective,” insofar as it holds the normative debate “hostage to issues in the philosophy of language and metaphysics” that are themselves contentious and possibly incapable of being settled (cf. Stubblefield 2005, 73). Presumably the biggest of such

issues are the ongoing disputes in philosophy between causal and non-causal theories of reference and between descriptivist and non-descriptivist theories of meaning. Now I think that Mallon is on to something here that we had better respect: we don't want to get bogged down in questions in the philosophy of language when we're trying to sort out answers to questions about race. Call that piece of wisdom *Mallon's caution*. Given Mallon's caution, I want to grant for the sake of argument that the disputes over reference and meaning are intractable (although I am not confident that this is so) and suggest that, even given that premise, two ways of pursuing the semantic strategy bypass this potential hazard.

One strategy for finding meanings for racial terms is to *separately* utilize rival theories of meaning, in order to show that no matter which side of the linguistic fence one is on, each independently requires us to adopt a certain ontological position with respect to race. Appiah (1996, 32–74) follows this path in arguing that whether we adopt an “ideational” view of meaning, according to which a term's meaning is determined by what people think about the term, or a “referential” view, according to which a term's meaning is determined by the nature of that to which it applies when we use it, we will end up with a set of biologically oriented meanings for racial terms that have no biologically interesting referent in the world. Realists could in principle adopt an analogous strategy.

Mallon's (2006, 549) concern about this first strategy is that “there is no reason to believe that all the plausible [theories of meaning]⁸ converge on a single answer regarding whether or how race exists,” as evidenced by the fact that so many different semantic premises have been marshaled to support different ontological conclusions. However, there is another way of evaluating this evidence: rather than it providing a reason to stop the semantic and therefore ontological parts of the race debate, we should take it as a reason to do more philosophy. That is, the appropriate response to semantic puzzlement is not to give up on conceptual analysis, but rather to do more analytic work to resolve the puzzlement. As such, whether different theories of meaning can independently converge on a univocal semantics of race that supports a single ontological conclusion is what the arguments are for on this first viable, cautious way of pursuing the semantic strategy: Appiah presents arguments concluding that this is exactly the case, and Paul Taylor's (2000) response is based in part on demonstrating that Appiah's analysis is faulty. For current purposes, the point is not to crown a champ in this dispute. Rather, it is to make

8 Mallon here actually talks about theories of reference, rather than theories of meaning, but as Appiah (Mallon's foil) is concerned with different theories of meaning, the latter seems to be the more relevant domain. In any case, the point should hold whether we are focused on reference or on meaning.

plain that the conversation is an important one to have, and that there is no reason to expect that some arguments cannot be given to help us come up with some answers.

The second way of identifying the meaning of racial terms without being held hostage to developments in the theory of meaning and reference is to analyze racial concepts in a manner that is, from the get-go, independent of debates among the rival theories of meaning and reference. A given term's referent (to focus on the theory of reference) is normally underdetermined by contentious theories of reference. Theories of reference are designed to account for various pre-theoretical semantic intuitions, such as that 'water' refers to H₂O, and if we can identify those intuitions independently of a theory of reference, as it seems we often can—after all, it is because of those intuitions' pre-theoretical appeal that the rival theories of reference can use them as evidence for their views—then we don't need a theory of reference to determine those terms' referents. For instance, Bill Clinton knows who the name 'Al Gore' refers to without, presumably, being able to spell out the correct theory of reference. And, of course, most of us can similarly know the referents of many of our terms. (This is not to deny, though, that there may be some hard cases where it is difficult to choose a referent without a settled theory of reference, nor even that there are some indeterminate cases, where a term has no stable referent.)

So, if we have an independently plausible definition of terms like 'race,' we can heed Mallon's caution by making sure that the proposed definition is consistent with plausible theories of reference and meaning, perhaps by stopping the search for a definition at a theoretically superficial level—or even at the level of analyticity—so that it hovers above debates over reference and meaning. And we see such attempts in the literature. For instance, Michael O. Hardimon (2003) provides an intuitively plausible analysis of the ordinary concept of race (which will receive extended treatment in the next chapter) that does not rely on any controversial theory of reference or meaning.⁹ Now, this argument makes an inductive leap: just because we have been able to analyze racial terms without appealing to a deeper theory of meaning and reference, this doesn't decisively prove that, in the case of racial concepts in particular, we will *never* get stuck in a referential jam that requires some more controversial tools. However, the evidence for the inductive leap is potent, so before entering the semantic debate we should be confident, if not entirely certain, that we can analyze

9 Slightly more technically, I think that Hardimon's analysis could be worded in the language of either descriptivism or direct reference theory, to either present a description of racial groups or to provide a way of talking about the racial groups directly referred to.

racial concepts without being held hostage to disagreements in the theory of reference or meaning.¹⁰

For these reasons, then, we should not endorse Mallon's (2006, 550) claim that "[i]f the only source of disagreement about 'race' talk were semantic, we could simply pack up and go home." It seems to me that if this were the only source of disagreement, the work would fall directly into the lap of those whose job it is to do conceptual analysis. For that matter, even if conceptual analysis turns out to be a fool's errand, that meta-philosophical question is itself something to be settled by doing philosophy. (Again, though, we should allow that some non-philosophers, including psychologists and cognitive scientists, also have some important contributions to make to conceptual questions about race, among other domains of inquiry.)

In the end, the exclusionists urge philosophers to focus exclusively on normative issues and forsake the ontological and conceptual discussions, and Mallon (2006, 551) in particular asks us to follow Sally Haslanger's (2000) recommendation to attend to what racial discourse *should* be like. I want to close the case for inclusionism by recalling another lesson from Haslanger: we need not choose between these two enterprises. The ontological and conceptual debates need not take place at the expense of tackling the undeniably paramount normative questions that many in the race debate are concerned to address. Indeed, the projects are complementary, for we'd be better informed as to both what our racial discourse *should* look like and how best to effect any necessary changes if we knew the truth about what it actually *does* look like. It is to this question that I now turn.

10 Note that Hardimon's theory-of-reference neutrality is but one instance of another philosophically ubiquitous method: to note a couple of conspicuous examples, Gettier cases are supposed to inform our analysis of knowledge, and Frankfurt-style cases are supposed to inform our analysis of blame- and praiseworthiness, but neither presupposes any particular theory of reference or meaning.