

## Epithets and Attitudes

A word is a slur when it is a conventional means to express strong negative attitudes towards members of a group, attitudes in some sense grounded in nothing more than membership in the group. A slur on Asians, for example, is a word which speakers know (and as competent speakers are expected to know) is used to insult and display contempt for Asians merely because that is what they are. What makes a word a slur is that it is used to *do* certain things, that it has (in Austinian jargon) a certain illocutionary potential. Given what slurs are used to do, it is no surprise that their use often achieves extreme effects on their targets—humiliation, subjugation, shame.

Slurs can be used without displaying contempt or causing hurt.<sup>1</sup> This happens, for example, when a slur is appropriated by its targets: it is an insult to no one, save perhaps the homophobe, for gay people to call themselves queer. A slur can be self-ascribed to record one's status as a victim of discrimination or worse. There need be no racism in an epithet's use by comedians to make fun of or criticize various attitudes and behaviors of both he who slurs and he who is slurred. One may use a slur in order to teach someone that it is a word which shouldn't be used. And an epithet can sometimes be used non-offensively in indirect discourse or narrative to portray someone else's racist remark or attitude.<sup>2</sup>

I will for the most part concentrate on uses of slurs which are offensive. My primary interest is the relation between such words' illocutionary and perlocutionary properties—their potentials for performing acts and achieving effects—and their more straightforwardly semantic properties; in particular, their potential for contributing to the truth and falsity of what a sentence says.

It seems undeniable that racial and ethnic slurs have application conditions. If I point at Prince Charles and say 'He's a Frog', I have—over and above any

<sup>1</sup> I am here indebted to conversations with Nancy Bauer and Margaret Klenck, comments from Robert May, and to the discussion in Randall Kennedy (2002).

<sup>2</sup> Strikingly it is not always possible for us, even when we want to, to use a slur in these ways without being subject to censure. No matter how honorable my intentions, I cannot join in the appropriation of slurs on African-Americans—at least not without something very much like an invitation from a target of the slur to so use it. (A nice discussion and example of this is given at the end of chapter 1 of Kennedy 2002.) Reports of slurring speech in which the slur is used are in some situations no less offensive than direct use of the slur, no matter how benign the reporter's intention.

moral breach—made a linguistic mistake, one of the same sort as I make if I point at him and say 'He's French'. Linguistic competence requires knowing that the French, not the English, are called 'Frogs', the English, not the French, are called 'Limeys'.<sup>3</sup>

Normally, if it makes sense to speak of a predicate being misapplied, and there are cases in which the predicate is not misapplied, there are going to be simple sentences in which the predicate is applied to an object which are true, and other such sentences which are false. At first blush, this would seem to be true of slurs and epithets: 'Frog' is a derogatory term for the French, and a derogatory term for the French is a term for the French. But if S is a term for the French and f is in fact French, then if I point at f and utter *He is an S*, I speak truly.

Appearances notwithstanding, one has a strong intuition that this can't be right. Let S be some odious racial slur. Imagine standing next to someone who uses S as a slur. Perhaps you are in front of a building where targets of the slur live or work; the racist mutters *That building is full of Ss*. Many of us are going to resist allowing that what the racist said was true. After all, if we admit its truth, we must believe that it is true that the building is full of Ss. And if we think that, we think that the building is full of Ss. We think, that is, what and as the racist thinks. This certainly seems to make us complicit in the racist's racist attitude, and thus to some extent racists ourselves.

One hears it said that the racist's utterance is true, but objectionably couched.<sup>4</sup> I don't think invoking the idea of unacceptably expressed truth helps here, for we cannot draw the right sort of line between thinking that it is true that S and thinking that S. The thoughts you think true, whether you like or not, are *your* thoughts. If you understand the racist and think what he says is true, you agree with him—you think (say) that the house is full of Ss. And we just saw where that leads: if you think this, you think of the people in the house as the racist does when he slurs. But if you think of people as the racist does in slurring, you are being racist.<sup>5</sup> "*The line between thinking it is true that S and thinking that S.*"

In what follows, I discuss the intuition that we cannot ascribe truth to utterances such as the racist's. I first consider some attempts to bolster the

<sup>3</sup> Of course the application conditions of epithets are vague; in the case of racial epithets, the use of the terms is arguably embedded in a mistaken view that races are 'real kinds'. Neither of these facts casts doubt on the claim that understanding such terms requires recognizing that their use is subject to correction for misapplication. Vague terms have application conditions, even if those conditions are vague. And the fact that, for example, jade is not a true kind doesn't mean that we cannot speak of right and wrong applications of 'jade'.

<sup>4</sup> Alan Gibbard, for example, has said this. Such views are discussed in Sections 1.3 through 1.6.

<sup>5</sup> There are lines one can plausibly draw here. On many views of thought, the thought that S is distinct from the thought that it's true that S, since only the latter must involve the notion of truth. And so, Donald Davidson notwithstanding, one may think that S without thinking that it's true that S, simply because one may have thoughts without having the concept of truth. One may plausibly say that someone who knows what it is for a thought to be true could think that S without thinking that it's true that S, because one may have beliefs in an area of discourse one thinks isn't truth-apt.

intuition with tools from the philosopher of language's tool kit, by saying (for example) that slurs have application conditions that nothing can satisfy, or saying that a slur introduces a faulty presupposition. No such account, so far as I can see, is terribly plausible.

I then sketch what I think we should say. As I see it, to think or talk slurringly of a person is, among other things, to have certain attitudes towards him, including evaluating him negatively and having contempt for him because one takes him to be of a certain race, ethnicity, religion, etc. The difference between thinking that Prince Charles is English and thinking that he is a Limey is, in part, that one is contemptuous of him when one thinks him a Limey, and thus *thinks* of him negatively when one thinks him a Limey. The attitude—the contempt—is *part of what one thinks*. Furthermore, to have an attitude of contempt towards someone because of their race or ethnicity is, *inter alia*, to represent one's target in a certain way: as contemptible because of his race or ethnicity. Such a representation is incorrect: no one is contemptible for such a reason. So what one says cannot be true. But the right attitude to take towards someone who slurs another is not that they have made a mistake that renders their thought false. Rather, we should reject the very way of thinking the thinker used in his thought. Not all representation is aptly evaluated in terms of truth and falsity. So, at any rate, I shall argue.

Slurs are 'thick terms', terms whose use, as Bernard Williams put it, 'mixes categorization and attitude'. Of such a term one wants to know exactly how (if at all) the attitude contributes to the thought expressed by using the term, whether the term could be understood by someone lacking the attitude, and under what conditions (if any) the presence of the attitude blocks the thoughts expressed by using the term from being true or false. I take up some of these questions at the chapter's end, arguing that in some ways everybody who has weighed in on these issues has gotten things wrong. Those (like Williams and John McDowell) who think that certain attitudes are in some way essential to some thick concepts are wrong; but so are those who say the attitudes annexed to a thick term neither help individuate the thoughts expressed with the term nor contribute to how the term represents. As I see it, whether two people express the same thought with a sentence is something which turns as much upon the interests of and social pressures upon the person *asking* whether the thoughts are the same, as it turns on facts independent of the situation in which we ask whether what one person

But while we can deny that anyone who thinks that the Frogs are a haughty lot must also think that it's true that they are, can we deny the converse? It is not clear to me how one goes about doing this.

Neither, it seems to me, is there much hope in suggesting that I might know what your utterance says, be so related to it that I could say what you say by using the very words you do, think that what you say with your words is true, but not think the thought in question because I refuse to 'mentally token' (or audibly token) your words. I'm sorry: If you understand me when I utter 'George Bush is a goddamn demagogue' and think to yourself 'How true, how true', you too think that Bush is a goddamn demagogue, whether you are willing to blaspheme or not.

says or thinks is the same as what another person does. For that matter, whether a thought is true or false can be as much a matter of the context in which it is asked *Is that true?* as it is of how things are independently of the asking. So, at any rate, shall I argue.

## 1.1

A slur and its 'neutral counterpart' have, as we might put it, the same targets—they are applied to exactly the same objects. But perhaps the slur is not *true* of targets.

How could this be? The meaning of a slur, like that of any word, depends on how it is used. A slur is a device which is used to express contempt for, to deride, and to insult its targets. It is mutual knowledge among speakers that slurs have such a use. Now while we expect that the nature and details of the negative attitudes expressed vary across the users of slurs, there is arguably rough uniformity in attitudes among those users. And so there is arguably a relatively small number of things normally conveyed by a slur's use, from which anyone who understands the slur will tend to draw its interpretation. I venture that most adults if they were asked about the attitude of a person who slurs African-Americans would come up with pretty much the same simple list: such a person despises or hates African-Americans because of their race, or believes that they are inferior because of their race.

This, one might propose, imbues slurs with a descriptive meaning beyond whatever descriptive meaning they have in virtue of applying to their target class. When the use of a slur is a conventional expression of an attitude *A* towards an object, held because the object is *F*, the slur comes to mean something like *worthy of A because F*. Because it has this sort of connection with the expression of racial hostility, a slur directed towards African-Americans means, to a first, rough approximation, something like *black and despicable because of it*.<sup>6</sup> But of course no one is or even could be despicable because of their race or ethnicity. So, the partisan of this view concludes, while the slur has application conditions, nothing does or could satisfy them. The same sort of thing is true of other epithets used to belittle on the basis of race, ethnicity, religion, sexual orientation, etc. Call such a view *DS*, for *defective satisfaction conditions*.

One might object to *DS* that it is possible to use a slur without insulting, demeaning, or evincing a negative attitude towards its targets. This happens, for example, when the slur is appropriated by its targets. If the slur's meaning is

<sup>6</sup> It is no objection to such a view, in my opinion, to suggest that shades of meaning are being obliterated here, as Jennifer Hornsby (2001) seems to suggest of a kindred story. For one thing, the view gestured at in the text can accommodate differences in application conditions so long as such differences are underwritten by differences in mutual expectations.

preserved in the appropriation—in particular, if there is no change in meaning relevant to truth-conditions—the fact that the targets of a slur can and often do find ways to defuse it by adopting it seems to show that DS is wrong.

The objection is inconclusive, for it is not at all clear that, for example, 'queer' preserved its meaning upon appropriation by the gay community. Before appropriation, it was arguably *conventionally* used to express hostility and homophobia: a good dictionary would have marked it as derogatory. After appropriation, it was not—at least not by the appropriators—used to express hostility or homophobia. A good dictionary will today note that the word has a use which is not derogatory. There is a case to be made that in appropriation there was a meaning change.

In any event, the objection misses the point of the original worry about slurs and truth. Suppose illocutionary facts do not enter into meaning: that a word is used, even conventionally used, to insult or denigrate, let us momentarily suppose, is not a fact about a word's meaning, but 'simply about its use'. Suppose that the same is true of whatever 'expressive' properties a word might have, so that the fact that a word is used to exhibit or otherwise to give vent to an emotion or evaluation is 'pragmatic, not semantic'. (I do not think this is a sensible way to use 'meaning', but let this be granted for a paragraph.) If we grant this, we are well on the way to agreeing with the objector that there is nothing more to the meaning of a slur beyond what is common to the use of the bigot and the use of the appropriator. But one can *still* argue that when someone utters a sentence with a slur and intends thereby to insult, denigrate, or devalue on the basis of group membership, what the sentence says is not truth-apt. For at issue is whether the illocutionary or the expressive have the power to trump the semantic, so that the fact that a sentence is used in a certain way changes the terms on which it can be evaluated. If we decide to slice meaning so that the illocutionary is not part of meaning—so that it is 'merely pragmatic'—we need to acknowledge the possibility that there might be a determinant of truth and falsity beyond what a sentence means and 'the way the world is'. To object to DS from the possibility of using a slur without its standard performative or evaluative baggage misses this point.

DS depends upon a sort of sociolinguistic hypothesis, that speakers expect that they can read the attitude of a slur's user simply off the fact that he uses it. That hypothesis might be true of some epochs but not of others. My limited sociolinguistic take is that in my country's recent past DS may *not* have been true of a great many slurs, including some pretty odious ones. Forty or so years ago it was not uncommon to hear people use slurs on the playground and at the dinner table. People—including people one would not have taken to have dislike or contempt for the targets of a slur—apparently thought little or nothing of telling jokes in which such slurs occurred. They would use them to talk about their target classes. It was not, I think, implausible that many users of slurs were not consciously invested in or committed to hatred, moral condemnation, or

belief in the inferiority of the targets of slurs. Rather, the use of slurs reflected a complicated web of attitudes, including discomfort about or fear of what seemed pronounced physical and cultural differences, as well as a lamentable tendency to talk as others did.

I think people generally recognized that this was so—that is, they did not assume it was clear what attitudes to read off the fact (but only the fact) that someone used a racial or ethnic slur in conversation. I am *not* denying that many users of slurs were invested in hatred and contempt. Rather, I am questioning whether those who slurred were routinely conscious of such attitudes and whether there was a widespread presumption that whoever used slurs must have had such attitudes.<sup>7</sup>

So I'm suspicious of DS as a *general* explanation of our reluctance to ascribe truth to slurring speech. And there are other reasons to question it. For one thing, DS conjures satisfaction conditions out of illocutionary action in a somewhat puzzling way. Someone addressing an Anglo audience who uses a slur on Asians typically thereby denigrates and insults Asians. So much is clear. But why should we think that he *asserts* that they are inferior or worthy of contempt? After all, do I *assert* that you are a jerk, a loser, or worthy of contempt if I stick out my tongue or give you the finger? Do I, for that matter, *assert* that you are a jerk, a loser, or worthy of contempt if in frustration I yell at you 'Jerk!!! Loser!!!!'? Do I *assert* that a certain bicycle chain is an object of my frustration, or deserves to be cursed, if, after a half hour of being unable to tighten it properly, I remark 'BLANKETY BLANK PIECE OF BLANK'? Of course not. Why should we say anything different about slurs? What is obnoxious about a slur is what is typically *done* with it.

Furthermore, DS doesn't explain one of the most striking things about slurs, something they share with other devices of invective and insult. Slurs typically 'scope out': their use is insulting and denigrating even when they are embedded. If \* is a slur on Jews, then uttering 'Bob didn't marry a \*; he married a Baptist' is as offensive, and offensive for the same reasons, as is uttering 'Bob married a \*': each sentence signals a negative evaluation of Jews. Almost any embedding of a sentence in which a slur is used is offensive; someone accused of anti-Semitism for uttering 'If Bob married a \*, then so did Ted' cannot plead that he only said 'if'. It's not clear why this should be if the odium associated with such slurs is located in their satisfaction conditions. Indeed, if that was all there was to be said about the terms, then it should be perfectly all right to endorse such things as the last-mentioned conditional, for such sentences would be non-problematically true, having antecedents which couldn't be.<sup>8</sup>

<sup>7</sup> For some relevant discussion in the philosophical literature see Appiah (1990) and Garcia (1996). There are a number of books containing sometimes lengthy transcripts of conversations with 'ordinary people' from this era concerning their racial attitudes which give some sense of how easily slurs were used; see, e.g., Wellman (1977).

<sup>8</sup> Someone might suggest that we can explain our reticence to ascribe truth to claims made with slurs by supposing that the satisfaction conditions of slurs are vague. If slurs have vague satisfaction

## 1.2

In this regard slurs resemble expressions like descriptions and clefts. The use of a definite description *the F* signals an assumption that there is a unique F. The assumption 'scopes out', insofar as the description's use signals this assumption even when the description is inside a negation, a conditional, or a question. Likewise, for the cleft construction *It was X who F-ed* and the assumption that someone F'ed.

Descriptions and clefts not only signal speaker assumptions, they introduce conversational presuppositions. If I say 'It was Cheney who plotted to steal the oil', I signal my assumption that someone plotted to steal the oil. If no

conditions, presumably this is because, given a slur S, while it's true that S's satisfaction conditions are A or they are B or . . . or they are Z (for more or less non-vague conditions A through Z), it's not determinate that S's satisfaction conditions are A, and not determinate that they are B, and . . . not determinate that they are Z. One way to spell this out makes the indeterminacy one among conditions false of a slur's targets. So it would go if someone said that (a) it is indeterminate whether a slur on target T meant *inferior because of ethnicity T*, indeterminate whether the slur meant *worthy of contempt because X, Y, Z* (X, Y, Z being stereotypes of target T), indeterminate whether it meant *less pleasant to be around because of ethnicity T*, but (b) it is determinate that the slur means *inferior because of ethnicity T*, or it means *worthy of contempt because X, Y, Z*, or it means *less pleasant to be around because of ethnicity T*.

This view is liable to the two objections at the end of this section. The view is conjuring satisfaction conditions out of illocutionary acts just as much as more straightforward views which assign determinate satisfaction conditions to slurs; the view makes mysterious the fact that we find the denial of a claim made with a slur just as offensive as the claim itself. (For if it is determinate that sentence S means A or means B, and A and B are both definitely false, surely S is definitely false, too.)

Someone might say that it is not just vague *what* the satisfaction-conditions of slurs are, but vague *whether* anything satisfies them. If it is vague whether 'queer' as used by the homophobe means *homosexual* or means *contemptible because homosexual*, such a view would be true of 'queer'. This would presumably make a use of *He is queer* applied to a gay person truth-valueless (because it would be indeterminate whether it means the true *He is homosexual* or the false *He is contemptible because homosexual*), and it would make such a use of *He is not queer* also truth-valueless (this being indeterminate in meaning between the false *He is not homosexual* and the true *He is not contemptible because homosexual*).

There are at least three reasons to be unhappy with such a view. First of all, it leaves us with no explanation of why we resist ascribing truth to a wide range of claims made with slurs. One's reaction to the homophobe who says 'Sean isn't queer, but I think maybe Bob is' does not depend on the sexual orientation of Sean and Bob—but on the current view what is said would be true if Sean were not gay, but the speaker believed that perhaps Bob was. There is a use of 'real man' on which it (seems to) mean something like 'not gay'; one wants to say that the problem with 'Real men aren't queers' is the same problem we find in a sentence like 'He is queer'. But the former—at least on natural accounts of truth of vague talk—would turn out true on the current view. The second problem with this view is that it continues to conjure satisfaction conditions out of illocution. The third, and I think primary, problem with this view is that it misdiagnoses the (probably) genuine vagueness of slurs. It probably is vague what 'queer' means in the homophobe's mouth—there are any number of things which bring the homophobe to fear and dislike the gay person, and settling on one of these as 'the meaning of "queer"' obscures this. But the idea that the homophobe 'sort of means' nothing more with the term than 'homosexual' is surely wrong: the term, as the homophobe uses it, is *definitely* derogatory.

one objects, the assumption 'becomes part of the conversational record'. This means that subsequent conversation will be governed by the assumption that someone so plotted: I and my audience may use the claim as a premise in argument, a counterexample to others' assertions, etc. Perhaps slurs—at least in their core uses—not only signal negative assumptions on the part of the speaker, but—when no one objects to their use—introduce a negative presupposition about the slur's target into the conversational record.<sup>9</sup>

That the use of a slur signals the *speaker's* belief or assumption that a particular group deserves negative evaluation isn't enough to explain why we are loath to ascribe truth to a sentence in which it is used. A speaker's uttering 'I think Eastern Europeans are contemptible' signals such an assumption, but I will accept the utterance if I think the utterer sincere. But if a slur not only signals a speaker's assumption but makes it a conversational presupposition if no one objects, that explains the intuition with which we began. For when a presupposition we reject accompanies a sentence's use, we resist calling the sentence true. And if a slur presupposes that its target is contemptible because of race or ethnicity, it presupposes falsely.<sup>10</sup>

What is it for something to be presupposed? Let's follow Robert Stalnaker, and say that someone in a conversation presupposes claim p provided he assumes or believes—or at least is disposed to behave as if he assumes or believes—all the following: p; that other conversants assume or believe p; that conversants recognize that he is making these assumptions, or has these beliefs.<sup>11</sup> A conversational presupposition is something all the members of a conversation presuppose.

<sup>9</sup> If this is correct, slurs are *really* pernicious devices. For when all the members of a conversation presuppose p, it typically becomes mutual knowledge in the conversation that all presuppose it. So if my not objecting to your using a slur inscribes on the conversational record that its targets are contemptible, it inscribes on the record that I assume that they are contemptible. If slurs are devices for introducing presuppositions in the way that descriptions and clefts are, they are devices which through our silence make us complicit in the bigotry of others.

<sup>10</sup> Some say that when a sentence's use has a false presupposition, the use is neither true nor false. Others say that presupposition failure in and of itself does not cause a sentence use to lack truth-value. Since my conclusion will be that slurs don't trigger presuppositions to begin with, there is not much point in our wading into the waters of this dispute.

<sup>11</sup> Stalnaker (1974: 49–53). (The definition in the text is adapted from Stalnaker's definition of what it is for a *speaker* to presuppose p; the adaptation consists in replacing 'addressees' in Stalnaker's definition with 'conversants'.) The invocation of dispositions is to handle cases in which (for example) a speaker uses a description ('I can't be at the meeting; I have to pick up my wife's sister at the airport') realizing that the audience may not know or assume that it denotes, but expecting this to become part of the conversational record once he uses it. This sort of use is a central case of presupposition; Stalnaker suggests it can be captured by identifying presupposition with a *disposition* to behave in conversation as if the relevant beliefs or assumptions are present.

I myself am inclined to analyze these cases differently from Stalnaker. I would say that I do not presuppose *in the sense of presuppose for the purposes of the conversation* that I have a sister when I say out of the blue 'I must get my sister', but I will presuppose it after I have said it, provided no one objects; I will, in that case, also expect you to as well. If we say this about such cases, we can leave presupposition as a disjunctive attitude (belief or assumption), instead of identifying it with a disposition. This is relevant to the issues which n. 12 raises. Since these issues are something of a sidebar, I do not explore them here.

We should linger on the notion of assumption. Assumptions are assumptions for particular purposes, and 'assume' here is elliptical for something like 'assume for the purposes of a conversation'. We will assume or presuppose all sorts of things we do not believe, or even that we disbelieve. So it goes in proof by *reductio*, or in a conversation in which we humor someone's delusions. To assume *p* for conversational purposes doesn't require much, if anything, beyond being willing to let go remarks made by others which entail *p*, to allow *p*'s use as a premise, and so on.

What we come to assume for conversational purposes we often assume automatically, without conscious fanfare. One just *does* assume (say) that there is a positive solution to the equation on p. 331, having read 'suppose the equation on page 331 has a positive solution'. No conscious decision is made, no bell audible to consciousness rings. Given the bloodless nature of assumption (its not involving belief), its ability to fly below consciousness, and its more or less automatic nature, one can't dismiss out of hand the idea that a slur's use might routinely inspire a negative assumption about its target.<sup>12</sup>

The idea we are pursuing is that slurs introduce negative presuppositions about their target into the conversational record when no one dissents. There are two ways to spell it out. One might first of all say that this is in fact the typical result of a slur's use: normally, an unchallenged use of a slur is followed by auditors presupposing something negative about the slur's targets. Call this hypothesis *DF* for *de facto*.

*DF* is surely false of contemporary uses of slurs addressed to their targets, or used with the knowledge that a target is in the audience. There is little reason that I can see to think that the target of a slur typically shares, or can be made to assume for conversational purposes, the negative attitude of someone bigoted towards him. This is not to say people are not sometimes bullied into 'accepting' an insulting remark or its presuppositions, nor that it isn't sometimes the wisest course for a slur's target to 'let it pass'. Rather, it's to say that a slur's target will not normally make negative assumptions about himself when slurred, and so *DF* is not even roughly true of 'second-person' uses of slurs.<sup>13</sup>

<sup>12</sup> Does the fact that I don't believe and could not be brought to believe that Hispanics are contemptible speak against the idea that when someone uses a slur on Hispanics in my presence I automatically assume that Hispanics are contemptible? Certainly we would need some argument that this is so. Assumption, as I have been trying to point out, is something which occurs in a mental area where non-deliberate behavior is rife. It is not clear that one's moral beliefs will have the power to prevent the sort of quasi-automatic effect which assumptions are. Perhaps they do, but if they do, this is a surprising fact for which we need evidence.

Relevant here, perhaps, is the fact that there is nothing wrong *per se* with assuming the claim that Hispanics are contemptible: If I assume this in order to perform a *reductio* on a racist's beliefs, I haven't done anything wrong. Whether an assumption is to be condemned turns in good part on *why* the assumption is made.

<sup>13</sup> Suppose we are in a historical era where bigots are in power and the mere sign of opposition on the part of the target will be dealt with harshly. Then most targets may well be disposed to behave as if they assumed that they were inferior when slurred, such behavior being the safest way

Is *DF* true of slurs not addressed to nor made in the presence of their targets? It depends upon the historical and cultural context. If prejudice is endemic, something like *DF* may well be true. But if we are looking to *DF* to explain *our* reluctance to ascribe truth to slurring talk, it must be true of current practice. This seems implausible to me. What can a speaker who uses a slur and is not challenged conclude, about whether he would be challenged if he were to start making derogatory assertions about the slur's target? This depends on the particular case. In *many* cases, people who might let a use of an epithet pass would not let the direct voicing of the attitudes underlying the epithet pass unchallenged. So the unchallenged use of a slur normally does not lead to our presupposing the user's attitudes towards the target.

A slur might not *in fact* introduce a negative presupposition, but still be something which is *supposed* to do so. A description of a slur's meaning would, on this view, mark it as a word which has this purpose, so that a competent speaker would know that if no one objects to a slur's use, the user and other conversants are entitled to assume that all presuppose something negative about the slur's targets. Call this hypothesis *DJ* for *de jure*.

I assume that most contemporary speakers don't want words with the properties that *DJ* ascribes to slurs. One wonders how there could be such words if most speakers don't want them. Suppose it is known that most people do not think group *G* contemptible, and known that people are not inhibited by fear from saying that they do not think this. Are you entitled to think that if you insult or belittle the members of *G* and I do not choose to call you on it, I am willing to let you do it again, that I would let you go on to voice claims about inferiority or the like, or that I share your contempt for *G*? Surely not. But then it is not clear how the sort of expectations necessary for *DJ* can be in place when prejudice is not widely shared.

As I see it, *DJ* misdiagnoses 'how slurs work'. Is it *the rule*, or at least a good rule of thumb, that someone who uses a slur is trying to slip an assumption into the conversational record? Surely not. A pretty good rule of thumb is that someone who is using these words is insulting and being hostile to their targets. But there is a rather large gap between doing that and putting something on the conversational record. If I yell 'Smuck!' at someone who cuts me off, or say 'That smuck who just cut me off should lose his license', I insult and evince hostility. Am I entitled to assume, if you don't say 'He's not a smuck', that you assume that

to behave. It seems to follow that in such a situation the use of a slur in the presence of a target would be followed by the target's presupposing that he was inferior.

This may seem odd, but I think it is the right result. To assume *q* for purpose *P* is to be prepared to allow *q* to play a particular role, that role depending upon *P*. For example, to assume that Iraq has a stockpile of WMDs for the purposes of charting public policy is to be prepared to take seriously reasoning about public policy whose premises imply that it is likely or certain that Iraq has stockpiled WMDs. To assume *X* for purposes of conversation is to be prepared to let people reason from *X* in conversation, not to challenge *X*, etc. And it is a fact that one *can* make people make pernicious assumptions for many purposes, conversational ones included, via intimidation.

the person in question is a smuck, or are hostile towards him? Surely not. One doesn't have a conversational obligation to demur, if one disagrees, from whatever slings and brickbats one's conversational partners performatively launch. (One may, of course, have some other obligation to demur.) In this regard, the 'content of an insult' is different from what is asserted and from a conventional presupposition. One *does* have a conversational obligation to demur from these if one disagrees. For if one does not, *they* go on the conversational record.

Slurs are not conventional means of inserting contemptuous attitudes into a conversation's record. They pattern with terms of insult and invective—'dweeb', 'smuck', and so on—when the latter are addressed to someone or used as insulting prefixes. If anything is conventional about a slur, it is that its third-person use gives vent to a malevolent attitude, its second-person use insults and denigrates.

## 1.3

Perhaps the intuition with which we began is wrong, and racial and ethnic slurs *are* true of their targets. Perhaps they are, as Alan Gibbard puts it, terms which combine

classification and attitude . . . where the local populations stem from different far parts of the world, classification by ancestry can be factual and descriptive, but, alas, the terms people use for this are often denigrating. Nonracists can recognize things people say as truths objectionably couched.<sup>14</sup>

I dismissed this idea at the start of this discussion. But perhaps it deserves another hearing if appeal to satisfaction conditions and presupposition failure are unable to explain how the sentences we have been considering can fail to be true.<sup>15</sup>

The idea of a concept which combines 'classification and attitude' is of course Bernard Williams's idea of a 'thick concept', a concept whose application is 'determined by what the world is like . . . [and] usually involves a certain valuation'.<sup>16</sup> Canonical examples are terms for virtues, vices, and sins. Since the

<sup>14</sup> Gibbard (2003a: 300). I should mention that what I have elided here is the claim 'Racial epithets may sometimes work this way'.

<sup>15</sup> At this point the philosopher of language is thinking—Aha, now comes the discussion of the idea that while utterance of *He is an S* will be true, if aimed at a target of the slur *S*, there is a 'pragmatic' explanation of why we reject such utterances. As the savvy philosopher of language knows, such explanations usually invoke some broadly Gricean mechanism, holding that (a) the utterance suggests or conveys (but doesn't 'strictly and literally say') some claim which (b) we take to be false, while (c) we are so intent on rejecting this claim that we don't differentiate rejecting it from rejecting what the sentence 'strictly and literally says'.

I will discuss this idea, but only at the end of the chapter. My reason for deferring discussion is that it is hard to see what is wrong with this explanation unless and until we have the right explanation—which the next sections attempt to sketch—on the table.

<sup>16</sup> Williams (1985: 129).

application of such terms is determined by what the world is like, Williams says, they can be true of what they are applied to. So, to fix on an example, when Augustine said that his early life was unchaste, he spoke truly, his command of the criteria of chastity being after all impeccable.

Gibbard thinks that Williams is wrong about concepts like chastity—they are not true or false of anything—but that he is right about the concepts expressed by racial slurs. What, exactly, is the difference? Gibbard tells us that '[i]f how people classify with a thick term depends on their attitudes and their attitudes are faulty for their situation, then they do not get it right or wrong with their classifications'.<sup>17</sup> According to Gibbard, poor Augustine's classification of his behavior as unchaste depended on his 'faulty attitude', and so he can't be credited with speaking the truth. What exactly it would be for a classification to depend on an attitude, faulty or otherwise, is something to which I will return below. For the moment, though, let us concentrate on the idea that when a slur on Jews is used by an anti-Semite, the attitude expressed is faulty, but the classification made does not depend on the adjoined attitude. The *classification*, after all, could be pulled off with the word 'Jew'. So, Gibbard thinks, sentences using the slur can be true though 'objectionably couched'.

Gibbard's picture seems to be something like this. There are certain things—classifications, let us call them—which people make when they speak or think. At least some of the time, classification can be separated from the way in which it is made: the anti-Semite and the rabbi make the *same* classification when one uses the slur and the other speaks of the Jews. Classification, however it is made, is necessary and sufficient for the question of truth to arise: when and only when a classification is made, do we have something which—issues of vagueness put to the side—is true or false.<sup>18</sup>

I think this is a bad picture. I take it that what is in the first instance true or false are the things we think and say, the things with which we are concerned when we remark that *X said that A*, or *Y thinks that B*. When we think or say something which is true (or false), we do indeed classify. But classification (in the sense of the last paragraph) is to be explained in terms of the concepts or 'ways of thinking' which we employ when we think and say things. It is *because* the anti-Semite thinks of someone in a particular way when he slurs him, *because* we think of a person in another way when we think of him as a Jew, and because these two *different ways* of thinking of a person are related in certain ways, that we are entitled to say that there is a

<sup>17</sup> Gibbard (2003a: 302).

<sup>18</sup> This is oversimplified. There is classification in the command *Take out the trash!*, but the command is not a truth bearer. What I mean to ascribe to Gibbard is the idea that simply to classify is—vagueness aside—to do something which is either true or false.

It sounds a bit odd, I suppose, to say that someone like Augustine doesn't manage to classify behavior when he calls it chaste. This is however, Gibbard's view: 'if we reject Augustine's concepts, we cannot coherently credit him with achieving classification' (Gibbard 2003a: 298).

sense of classifying in which the anti-Semite and we are making use of the same classification.<sup>19</sup>

There are obviously significant similarities between thinking of someone as a Jew and thinking of him in the slurring way the anti-Semite does. In an important sense, the application of each concept is guided by the world in the same way. There is a kind of mistake which one makes if one applies either way of thinking to a non-Jew; there is a kind of success, complimentary to the mistake, which one achieves if one applies either to a Jew. 'Classification' is a good name for the activity liable to such successes and failures. Classification so understood is a prerequisite for truth, since what involves no classification is no candidate for truth.

But why think, as Gibbard does, that classification is *sufficient* for truth or falsity? The classification itself is a bloodless thing, something that cannot be identified with thinking of an individual as a Jew, or as the slurring anti-Semite does, for the classification is something made when one thinks of someone in either of these ways. Whatever truth belongs to a classification is truth it inherited from the thought expressed in making it.<sup>20</sup> To point out that the anti-Semite is 'classifying' correctly—and remember that 'classify' here is a *technical* term, one which is not synonymous with everyday terms like 'think of' and 'conceptualize'—seems no more relevant to the question of whether he has said something true than does the fact that the anti-Semite has successfully targeted his prejudice.

## 1.4

But why *not* allow the thought expressed in uttering *He is a \** to be true, when the individual classified is a target of the slur \*? Well, to say that a thought is true is, *inter alia*, to approve of it, to endorse it: if I say your thought is true, I am saying that you are representing the world aright. To say that your thought is true is to endorse it *as* representing the world correctly. But when the anti-Semite thinks slurringly of Jews, he is not representing the world aright. We cannot approve of representing or thinking of people in the way in which he does, for we would approve of thinking of people as inferior merely because of their religion. It is wrong to represent anyone in that way.

<sup>19</sup> Why say these are different ways of thinking of someone? Because to think that Jews observe the sabbath on Saturday is not, automatically, to think that \*s do so. But if the two ways of thinking were the same, to think the one would be to think the other. Remarks below about individuating what is said are, of course, relevant to this argument.

<sup>20</sup> This way of speaking may make it sound as if the point depends upon reifying what is said. I don't think this is so. I believe the argument here is just as good (or just as bad . . .) if we begin by insisting that it is particular uses of sentences (and particular mental states) which are in the first instance true or false.

I just suggested that we cannot endorse the anti-Semite's slurring performance as true, as it involves misrepresentation. Do not infer that I therefore think we must pronounce it false. To say that the anti-Semite's claim, that \*s do not observe Lent, is false is to say that he made a particular kind of error, a kind of error which can be corrected *merely* by adjusting the way he classified or quantified, or by judicious use of negation. But this is not the case. If it were, we could fix the problem which prevents us from ascribing truth to

(B) If Bob married a \*, then so did Ted.

by correcting the speaker along the following lines: Well, that's false, for Bob did marry a \*, but Ted did not. The kind of error—the kind of misrepresentation—involved when one thinks of someone in slurring fashion is not the sort of error which can be corrected merely by reclassifying. Sometimes—this is one of those times—someone may represent the world but, because of the kind of representation employed, truth and falsity are simply the wrong terms in which to evaluate the representation. We do not reject (B) because it is false, but because it involves a way of thinking of Jews that is just not a way that one can think accurately about them. Negating (B) is not enough, for thinking (what's said by) (B) negated is *still* thinking of Jews in an odious, inaccurate way.

Some will find my way of looking at thought and representation curious: How can it be that a thought represents things as being *so* (things contain a group despicable because of ethnicity), and the (truth-functional) *denial* of that thought represents things as being the same way? But surely it's not at all mysterious that this sort of thing might happen. We all accept the idea that a claim and its negation may both involve commitment to a representation's being accurate—this is just the sort of thing that happens when a claim carries a presupposition, for a claim's presuppositions project to its negations. For the reasons given in Section 1.2, saying that the (mis)representation involved in thinking that \*s are F, or in thinking that \*s aren't F, is a presupposition (in the semanticist's somewhat technical sense of 'presupposition') seems wrong. Still, quite clearly, when one uses a slur, over or under a negation sign, one *is* thinking of the word's targets in a negative way.

Some will find my way of looking at thought and representation curious: How can it be that there is a *representation* that is in no (relevant) way vague, but is neither true nor false of its target? Given a representation—*Things are thus and so*—mustn't we always be able to ask, Are things like that, or are they not? And unless vagueness (or the problems attending liar sentences) intervenes, mustn't there be an answer to this? But the answer to this question determines whether the representation is true (as things are as it represents) or false (as they are not).

There are two things to say in response to this, though they are things that it will take most of this book to spell out.

First, we can agree that whenever we have representation we can ask whether the representation is accurate and expect that there is an answer. We can do

this even when the representation is vague. We can say of the borderline case of baldness that he is not bald (denying the claim that he is bald) and that he is not not bald either (denying the negation of this claim). To say that we can coherently deny a claim and its negation presupposes that denial is a *sui generis* speech act, not to be identified with the assertion of a negation (or any other assertion), the denial of a claim being appropriate when the claim is not true—when it is false or without truth-value. As we will see in the next chapter, there are compelling reasons for thinking of denial in this way. Given that we can deny both *p* and its negation when *p* is without truth-value, this is just what we should do with the representations involved when one person says that all \*s are Fs, and his companion says, no, at least one \* isn't F. We should reject both of these claims, just as we should reject both the claim that Jo is bald, and the claim that Jo is not bald, when Jo is a borderline case of baldness.

The second thing to be said here is that there are good reasons, quite independent of our current worries about epithets, to think that some representations are not to be evaluated in terms of truth and falsity. As I will argue in Chapter 5, the natural way of thinking about claims, that various activities are cool, people attractive, or foods disgusting, is to understand them as not being candidates for truth or falsity, not even relative truth or falsity. We may just disagree about whether Johnny Depp is good-looking: you think he is, I think he's not. In such a case, though we really disagree, neither of us need be making a mistake. But then, though I *represent* Depp as not attractive, and you *represent* him otherwise, our representations aren't true or false. They can't be, for if they were, at least one of us would be making a mistake.<sup>21</sup>

I will be accused of punning in the last few paragraphs on 'wrong' and 'incorrect'. It might be morally wrong, someone will say, for the anti-Semite to think or say that X is a \*; it does not follow that it is not true that he is a \*. There are two sorts of normativity at issue, moral and semantic. Now I agree that saying that it is morally wrong to think *p* is not saying that thinking *p* is misrepresenting the world (and thus not thinking anything true). These are different kinds of normativity and I am not (at least I think I am not) confusing them. What I am claiming is that to think of someone as the anti-Semite does *is* to misrepresent them in a way that deprives what is said of truth.

When the anti-Semite thinks of someone in his anti-Semitic way, he thinks in a way that expresses, that vents his negative attitude towards Jews, and thereby shows contempt for and denigrates them. To *do* these things is to misrepresent Jews. It is to misrepresent them not because one is using a word that means something like *contemptible because Jewish*. Rather, it is to misrepresent Jews because one is doing certain things—e.g. expressing negative attitudes and

<sup>21</sup> I stress that the material in the last two paragraphs is not intended as a full response to the worry that preceded them; it is instead an indication of how Chapters 2 through 5 will develop a view responsive to that worry.

contempt elicited by religion—the doing of which is one way to represent Jews as worthy of contempt. To have or display contempt for someone, to think badly of them by having such contempt, *is* to think of them, to represent them, as worthy of contempt.<sup>22</sup>

I imagine some will deny that there could be a genuinely *representational* fact whose explanation depends upon pointing to an evaluative or affective component in someone's way of thinking about an object or property. Whenever a thought is *about Z* or a concept is *true of Z*, it will be objected, there is an explanation of this that makes no use of the evaluative or the affective. And if so, it will be said, I am wrong to say that just to perform, to evaluate, or to have an affective reaction is thereby to represent.

I think this objection is just wrong. Suppose that a certain scenario—invading Iran, say, to root out its stockpiles of WMDs—causes me to be afraid every time it is mentioned. It is not that I think 'Ooooh, that's SCARY'. Rather, there is something about the mere mention of it that makes me afraid for myself and my children. It *could* be, of course, that I subconsciously cognize the proposition that the scenario is frightening. But there is no reason to suppose that this will often, or even typically, be the case. Does the pogonophobe reason his way to his phobia by thinking 'How do beards frighten me? Let me count the ways . . .'? Of course not.

Whatever the mechanism of my fear, that I have this reaction to the scenario in itself means that I *represent* it as frightening. Note, in support of this, that we typically treat such fear—as well as phobias—in the way we treat all representations: as subject to criticism (or congratulations) for irrationality (or for the recognition of genuine danger). The upshot, as I see it, is that there is nothing confused in thinking that the affective is responsible for representational properties. And I do not see why we should say anything different about the evaluative or performative. That someone thinks, evincing contempt, *\*s are different from us* means that the person *represents* the people he is thinking about as contemptible. To evaluate someone or to display value-laden emotion towards him *is* to represent him in a certain way.

<sup>22</sup> There are some subtle issues here concerning acts and attitudes. Suppose that Smith has a belief he expresses so:

\*s don't observe the sabbath on Sunday.

He may have this belief even when he is not expressing it or consciously entertaining it. When he is not expressing or entertaining the belief, he is not *displaying* contempt for anyone. But the belief, even when not being expressed or entertained, presumably involves the same way of thinking of Jews as it does when it is expressed. Is it, then, a mistake to say that the relevant way of thinking involves *expressing* contempt?

It is still true that Smith is contemptuous of Jews, even when he is asleep. The way he represents them (even while asleep) is a way of representing them which represents them as contemptible; to have a belief such as this is to have contempt for the Jews. So long as this is accepted, the points in the text are correct, even if we are loath to say that a way of thinking itself expresses an attitude.

## 1.5

I will be accused of a conceptual confusion. It is one thing, it will be said, to represent someone as being thus and so; it is a different thing to have or show contempt for him or to evaluate him, positively or negatively. The contempt or the evaluation may, of course, be based on a representation. I may evaluate you negatively or be contemptuous of you because I represent you as a Republican and disvalue all that is Republican. But the evaluation and contempt are distinct from the representation. Likewise, the *representation* involved in thinking of someone as does the anti-Semite has nothing to do with—that is, is completely separate from—whatever evaluation or performance may accompany it. The story I just told, however, crucially depends on the idea that to make an evaluation, to have or to express contempt, is (in part) to represent someone in a certain way. So, anyway, might someone object.

I want to concede something to this objection. Not that I am confused, conceptually or otherwise, but that when we consider a way of thinking which involves an evaluative outlook, it is usually—perhaps always—possible to separate the evaluation involved in the way of thinking from the classification involved in it. This is obvious for ways of thinking associated with slurs and epithets. Every slur, so far as I can tell, has or could have a ‘neutral counterpart’ which co-classifies but is free of the slur’s evaluative dimension. Presently I will argue that something similar is true of thick terms in general.

Of course the objector will not be satisfied with *this* concession. According to the objector, the anti-Semite’s negative attitude towards Jews is contingently conjoined with the way he thinks of Jews when he thinks of them slurringly. So the thought he has, when he thinks that \*s do not observe Lent, can be separated from the negative attitude. To my mind, it is an embarrassment for this more or less Humean picture that it is apparently committed to saying that there is ‘nothing negative’ in the way of thinking of Jews involved in the anti-Semite’s thought. For if the thought is independent of the negative attitude, it must be possible to think the thought without having the attitude.

There is something odd about the objection, that since affect and categorization are clearly separable in the case of slurs, affect does not help individuate the beliefs expressed with slurs. The conclusion just doesn’t follow from the premise. Why should the fact that an evaluation is a ‘separate existence’ from a categorization show that the two can’t go together to make up a single entity? Does the fact that one can think of Aristotle as the author of the *Metaphysics* without thinking of him as the author of the *Nicomachean Ethics* (and vice versa) show that the sense of the name ‘Aristotle’ for a particular individual could not be given by ‘the author of the *Metaphysics* and the author of the *Nicomachean Ethics*’?

All this raises the question, What makes it the case that a certain way of thinking, of an individual or an attribute, is constituted by a certain set of components? When one reads Frege one walks away with a picture of ways of thinking as tidy, sound-bite sized packets of information, neatly roped off from the mass of one’s beliefs: this man thinks of Aristotle as *the teacher of Alexander the Great who wrote the Metaphysics*, that woman thinks of him as *the unobservant Greek philosopher who had the temerity to write that women had fewer teeth than men*. Surely our mental landscape is not clearly segregated into discrete packets of sense. The information we take to be information about a single individual is (presumably) a vast jumble, not necessarily sorted in terms of importance or likelihood of accuracy. There are few, perhaps no, connections among ‘pieces of representation’ which would make this aspect of a representation of an individual connect analytically with that aspect. Furthermore, our reaction to and mental take on an individual or attribute will almost always mix together description, affect, wish, expectation, and so on.

We do, of course, think about, compare, and identify the beliefs of different people. And the way we do this is often best understood by individuating beliefs in terms of those small, sound-bite-sized ways of thinking to which I just alluded. But when we do this, *we* are imposing order on the jumble of the mental landscape. When we look at others and ask ourselves if they share our beliefs, we make—we must make—choices, dependent on our interests and situation, about whether another thinks of a particular individual as do we, and thus about when the mentality of the other contains a belief contained within our own.

There is, I think, great plausibility in the idea that our practices of identifying beliefs and other attitudes across people rely in part on such more or less one-off identifications of ways of thinking across thinkers.<sup>23</sup> And it is also plausible that when we ask what someone else said, or whether we agree with what they think, we often (but not invariably) individuate and identify in ways which effectively insert into sense—insert into the thought—the evaluations, conatus, and feelings we and others have towards objects and activities. My point is not theoretical or normative, but descriptive. It is made by thinking about cases.

Consider those two renowned experts on sin, Pope Paul VI and Simon Blackburn. Those who have been titillated by Blackburn’s recent celebration of carnality know that he will have none of the prudishness, shame, and haughty condemnation of sensuality which exemplifies the attitude of the author of *Of Human Life*. When he wrote that encyclical, the Pope was well aware of the sexual landscape of the sixties. He was appalled. He condemned the behavior of the time—he said it was lustful. Is this something with which Blackburn agrees? The Blackburn, who writes of orgasm

<sup>23</sup> I have argued elsewhere for this—e.g. in Richard (1990: ch. 3).

it is as close to ecstasy—to standing outside ourselves—as many of us get. . . it fills our mental horizon. . . This abandonment deserves more than a moment's attention. It is a good thing if the earth moves.<sup>24</sup>

Of course Blackburn does not think these things are lustful.

My description of the views of the professor and the prelate are, I think, perfectly natural. I expect that you will find the description natural even if you happen to know that the title of Blackburn's book is *Lust*, and that the book is a celebration of, and extended argument with the Christian tradition over, that for which it is named. And, as a matter of fact, if one looks at the way the Church and the don define the subject, it is pretty clear that they are talking about exactly the same thing. Blackburn takes lust to be

the enthusiastic desire, the desire that infuses the body, for sexual activity and its pleasures for their own sake<sup>25</sup>

The Church posts the following definition on the web:

Lust is disordered desire for or inordinate enjoyment of sexual pleasure. Sexual pleasure is morally disordered when sought for itself, isolated from its procreative and unitive purposes.<sup>26</sup>

Both Blackburn and the Church think that the behavior I alluded to before was lustful; what they differ on is whether this is a reason to condemn it.

Wait! Didn't I say before that we can say that Blackburn and Pope Paul VI disagreed about whether the activity was lustful? Indeed. We can say that. And we can say that they agreed that it was lustful—just not in the same breath. In some cases, for some purposes, we individuate concepts and beliefs so that they incorporate affective attitudes. In other cases, for other purposes, we individuate concepts and beliefs so that they do not incorporate the affective. When we individuate in the first way, we say that people whose evaluations differ as do Simon and Paul's disagree about whether a certain behavior is lustful. When we do not, we do not.

### 1.6

I have suggested that the identity of concepts is to an extent constructed by us. It is determined 'in the context of the observer' whether the evaluative attitude annexed to the racist's use of a slur is a 'part of' the concept expressed by the slur. To say this is *not* to say that it is a matter in which, in typical situations, we have a choice. It is, rather, to say that it is facts about our interests and

<sup>24</sup> Blackburn (2004: 24–5).

<sup>25</sup> *Ibid.* 19.

<sup>26</sup> <<http://www.Vatican.va/catechism>> sect. 2351.

relations—conversational, social, and so on—which control how we may or must individuate the concepts of others in thinking about their beliefs.

If I am right about this, one expects—I expect, at any rate—that given the right conversational or historical situation, it will be possible to detach the evaluative component from a thick concept. And if that is right, then one expects that no matter how closely a thick concept is currently associated with a certain evaluative attitude, it is in principle possible for someone to use the concept to categorize things without having the evaluative attitude towards what is classified. One can always, if you like, perform the sort of divorce Blackburn performs in wresting lust from Christian condemnation.

It has been said that the evaluative component of at least some thick concepts cannot be detached from them. Of thick concepts in general, Williams writes:

How we 'go on' from one application of a concept to another is a function of the kind of interest that the concept represents, and we should not assume that we could see how people 'go on' if we did not share the evaluative perspective in which [a thick] . . . concept has its point. An insightful observer can indeed come to understand and anticipate the use of the concept without actually sharing the values of the people who use it. . . . But in imaginatively anticipating. . . the observer also has to grasp imaginatively its evaluative point.<sup>27</sup>

Gibbard seems to concur—he thinks, for example, that Augustine's classifications 'depend on his attitude'. McDowell has on occasion suggested that many thick concepts apply to objects in virtue of those objects meriting a particular affective reaction. If this is so, then presumably creatures who do not have the relevant affective reaction to at least some of the right objects could not have the concept, somewhat as someone blind from birth could not have the concept of pumpkin orange.<sup>28</sup>

It is reasonable to think that if Williams et al. are right, I am wrong. For suppose that applying a concept to an object required taking a certain evaluative viewpoint. There would then be a strong case for saying that the relevant evaluation was essential to the concept. And this seems to be what Williams, Gibbard, and McDowell have in mind. They think, roughly put, that to think of something as chaste one must see it as good because it reflects a certain set of values. If this is so, then it is not 'up to us' as to whether a positive evaluation

<sup>27</sup> Williams (1985: 141–2).

<sup>28</sup> For McDowell's views see the essays in part II of McDowell (1998).

Some will object to McDowell and Williams that we may acquire concepts by being members of linguistic communities whose members already have the concepts, another's word being as good as his concept. I am not sure what either would say to this: perhaps they would draw a distinction between 'parasitic' and non-parasitic concept possession, with the former sort of concept possession being said to be of little interest; perhaps they would distinguish being in a position to echo others' words and being able to think their thoughts.

My reservations about Williams's and McDowell's views do not rely on the idea that concepts, like colds and other viral scourges, are spread orally.

should be annexed to a mental structure which we describe as a person's concept of chastity.

I have two comments. First of all, I think it is worth insisting on something which Williams at least seems to have quite clearly seen: *Any* ethical or moral perspective is liable to criticism, even very radical criticism, 'from the inside'. That is: No matter what coherent cluster of thick terms someone might begin with, it seems possible that one might (coherently) come to doubt that the evaluations associated with those terms were proper. Suppose that one could not acquire the concepts of lustfulness, modesty, chastity, and the like without occupying something like the Christian perspective with regard to sex. (I tend to doubt this, but let it pass for the second.) We can certainly imagine someone acquiring these concepts and then thinking that in applying them they were misevaluating. Imagine that Augustine underwent a second conversion, to a way of looking at the world on which the fact that an act manifested lust was not a reason to condemn it, that an act or trait evinced a dedication to not being sexually provocative was not a reason to praise it. Augustine being Augustine, we can imagine that he would try to bring as many of the benighted as he could into his new perspective. The obvious way to do this would be to try to convince them that they were making mistakes in evaluation: that the fact that an act was lustful was not, in and of itself, a reason to condemn it; the fact that an act was chaste was not, in and of itself, a reason to praise it. Surely Augustine could say this sincerely, meaning to and succeeding in contradicting what Catholics thought, when they thought that something's being lustful (or chaste) was a reason to condemn (or praise) it. And in doing this he would not be saying anything incoherent or self-contradictory: it just wouldn't be responsive to Augustine to insist that he was laboring under a 'conceptual confusion'.<sup>29</sup>

<sup>29</sup> I suppose I owe some account of what an 'evaluative perspective' is. As I understand the notion, an evaluative perspective is a complex of classificatory and evaluative dispositions, the sort of thing which is constituted by dispositions such as the dispositions to:

classify an act as unchaste if it is sexually provocative and such that a normal person would know that is sexually provocative;

classify an act as lustful if it makes manifest a desire for sexual pleasure ('for its own sake');

condemn an act if it is believed to be either unchaste or lustful.

(Actually, what I intend is to identify evaluative perspectives with functional states which are defined by 'Ramsifying' on (the names of the thick terms in) collections of disposition descriptions like those displayed in the text.) Such perspectives are, of course, vague affairs, and I assume that if we thought it necessary to flesh out the notion, we would do it in a way that allowed people to share an evaluative perspective when there was sufficient overlap between their dispositions to deploy thick terms. The evaluative perspectives we actually find occupied by humans will usually provide some necessary and some sufficient conditions for applying their thick terms; real examples will probably rarely if ever supply necessary and sufficient conditions. I shall assume, along with Williams, Gibbard, and McDowell, that we understand what it is to 'imaginatively participate' in such a perspective without actually having the dispositions which constitute it. There is, I suspect, no one way in which one might do this—one might feel moved in the way those who occupy the perspective are moved, but intellectually distance one's self from it; one might have a combination of dry intellectual appreciation for what the perspective requires along with a motivational echo of

I see the moral here as this: No (thick) concept is immune from being used by someone who has dropped its evaluative trappings, in the way Blackburn has abandoned disvaluing lustfulness and Augustine in our toy example has abandoned valuing chastity. And since *we* can describe what would be happening in such a situation, *we* can correctly identify a conceptual structure as realizing a certain thick concept *without* making it essential, for a structure's realizing the concept, that it involve any particular evaluative attitude. For we can say, as I just did, that someone like Augustine would *not* be laboring under a conceptual confusion in saying what he said about lust.

I promised two comments on the idea that a particular evaluation might be somehow essential to a particular thick concept. Here is the second. It seems to me that conceptual identity can be in part a matter of conceptual history and social relations. Whether, for example, a certain concept I now exercise is the same concept that Thomas Aquinas exercised three-quarters of a millennium ago is determined in part by contingent historical relations between his mental life and mine. It seems to me, furthermore, that sometimes a concept can be misapplied, even in what seem to be central cases, for a very long time before we come to see how it is to be correctly applied.

Now if one puts these facts together, it is puzzling why someone would hold the blanket (almost a priori) view that a thick concept's evaluative component is in some interesting sense essential to it. We can, after all, make dramatic mistakes in applying our concepts. We can begin by misapplying them, misled by demagogues, wishful thinking, or plain bad theory, and then come to get it more or less right. Or we can begin by applying it aright, and then get it wrong. Why should this be so only with regard to the 'descriptive' parts of our concepts? We think—I think, anyway—that Thomas and Paul just had it wrong about lust. I think that Augustine just had it wrong about chastity. If I generally choose not to use the word 'chaste'—well, that is in part because it's not common in these parts to use 'chaste' in a neutral manner. But there is all the distance in the world between constant conjunction and essentiality.

### 1.7

I say that when someone slurs, what he says is not true, not false. I do not say that the slurring representation is to be denied truth-value simply because it is 'expressive' or performative. I think that sincere utterances of 'I promise to meet you' or 'I find the defendant guilty as charged' are first and foremost performances. But I have no animus against the idea that they may be true for all

the perspective; perhaps one can occupy a position via simulation, somewhat as the virtuous actor tries to occupy the mind of the evil character.

that. Rather, I deny slurring talk truth and falsity because to ascribe it such we must represent as does he who slurs, and to so represent is to represent wrongly.

One can represent 'performatively' or 'evaluatively' with invective milder than a slur and say something true. Consider 'asshole'. It would be an interesting exercise, no doubt, to try to spell out the meaning(s) of this word without using invective; I shall not try. We know well enough what it is for someone to be an asshole—there are plenty of them around, and all of us are occasionally assholes ourselves. Suppose that Smith is an asshole and he is at the door. If I say, referring to Smith, 'That asshole is at the door', I display contempt for Smith by calling him an asshole. *That* does not prevent what I say from being true.

Both 'asshole' and slurs are devices used to display contempt for people in virtue of their possessing certain properties.<sup>30</sup> The difference between the terms is that there is nothing intrinsically misrepresenting about the reaction voiced by 'asshole': the way assholes behave merits contempt. To represent an asshole as an asshole is to represent things as they are.<sup>31</sup>

There are, of course, a variety of epithets, and a variety of sorts of uses of epithets. What we should say about a use of an epithet will vary from case to case. Suppose, for example, that Farmer Bob says, referring to his horse, 'I've had that nag there for ten years'. Suppose he has had the horse for ten years, and that the horse has done nothing to merit his contempt. Are we to say that Farmer Bob does not say something true?<sup>32</sup>

It depends. One might, after all, say—as my dictionary seems to say—that 'nag' has a use on which it is synonymous with 'old, inferior, or worthless horse'.<sup>33</sup> Then if Bob's comment is simply an assertion, with no performative element, it's false.<sup>34</sup> If Farmer Bob made jocular use of 'nag', as we make jocular use of 'son of a bitch' when we use it to refer to a friend, then the remark is true. In any case, there is reason to think that Bob's utterance is truth-valueless only if Bob is expressing (unfounded) contempt. And if he is, well then he *is* misrepresenting the horse, isn't he? If one thinks of the horse in the way he is thinking of it, one misrepresents it, in expressing unfounded contempt. And then the argument of previous sections seems to apply.

I say that when the bigot slurs, what he says is not true, not false. But is that always the case? When the bigot slurs and is sincere, he says something, he expresses something he believes. Suppose the bigot, suspecting that the person he sees in front of him is classified by the slur \*, says 'Well, I believe that he, Smith, is a \*'. He says he believes that Smith is a \*; he does indeed believe it. So isn't

<sup>30</sup> One could debate this. It is probably somewhat vague (indeed, somewhat variable across users and uses) what attitudes and evaluations are being expressed with 'asshole'. That doesn't effect the root point here.

<sup>31</sup> The last two paragraphs are in response to queries by Jim Higginbotham.

<sup>32</sup> Tim Williamson asked essentially this question.

<sup>33</sup> *Random House Dictionary of the English Language*, college edition.

<sup>34</sup> I assume that ten years old isn't old for a horse.

what he says true? But if it is, then it can't in general be that when someone slurs in saying something, what they say is not true. But if this is not in general so, why suppose it is in the cases discussed earlier in this chapter? The reasoning before was that what was said, since it employed slurring, involved a misrepresentation which deprived it of truth. Why wouldn't the same misrepresentation, and ensuing lack of truth, occur in the present case as well?<sup>35</sup>

Call the speaker *B*. We might ask three questions about *B*'s utterance of 'I believe that Smith is a \*':

1. Does *B* believe *that* [namely, what *B* says with 'Smith is a \*']?
2. Does *B* believe that Smith is a \*?
3. Does *B* speak truly, when he says 'I believe that Smith is a \*'?

I agreed that *B* believes what he says he believes; the answer to (1) is 'yes'. I have granted that *B* believes what he says; what *B* says is that Smith is a \*. I'm willing to infer that *B* believes that Smith is a \*, answering (2) affirmatively. But then, since *B* said that he believed that Smith is a \* and that's true, *B* spoke truly, right?

No. It is perfectly possible that when *we* say '*B* thinks that Smith is a \*' we speak truly, even if *B* does not speak truly when *he* says 'I believe that Smith is a \*'. For, as we observed at the beginning of this chapter, it *is* possible to use an epithet without slurring. Attitude reports appear to be one place where this is possible, even for those who are not in a position to publicly appropriate a slur's use. Given that *B* slurs and I do not, and that whether one slurs affects the identity of what is said, it follows that *B* and I do not express precisely the same thought with our uses of 'Smith is a \*' in this case. Likewise, it is possible that (niceties about the difference between *B*'s use of 'I' and our use of 'B' to the side) the thought *B* expresses with 'I think that Smith is a \*' is not precisely the thought that we express with '*B* thinks that Smith is a \*'. It is not as if in answering (1) and (2) affirmatively and (3) negatively we are committed to saying that some thought is true and not true.

There is an obvious difference between my thinking what I think and *B*'s saying and thinking what he thinks: I do not slur (and thereby misrepresent) Smith; *B* does. If we take the performative and expressive aspects of an utterance to enter into the individuation of what the utterance says—and the burden of the last sections has been that we often can, even *must*, do this—this means that we need not, must not, *identify* what *B* says, in his slurring performance, with what we think to ourselves when we homophonically translate it. So even if we are correct to think that *B* believes what he says—correct to think that *B* thinks that Smith is a \*—it doesn't follow that *B*'s slurring performance is something we need approve by giving it the dignity of 'truth objectionably couched'.

One has a strong intuition that *B* is expressing a belief that he has, when he utters 'I think that Smith is a \*', and thus that it must be true that *B* thinks that

<sup>35</sup> This objection comes from Scott Soames.

Smith is a \*. I think this intuition is veridical. What is wrong is the—admittedly quite natural—move from this intuition to the conclusion that when B speaks, he speaks truly. As I see it, we make the move because we have been taught, when thinking about matters of truth, to prise away any performative or expressive aspect of the utterance (indeed, of the thought) we are evaluating. But this, I think, is a mistake. If we could only regain our pre-Fregean semantic innocence, we would, I think, find it plainly incredible, that it is irrelevant to which way of thinking is associated with a word's use whether that use expresses contempt. If we accept that B and we are expressing different ways of thinking of Smith with \*, we should be ready to accept that the 'truth status' of a sentence in which \* occurs may differ, when we use it and when B does.<sup>36</sup>

I say that sentence uses in which the user slurs say nothing true or false. One might complain that the cost of saying this is too high. Shall we, it will be asked, say that slurs have no extensions, or shall we instead say that the truism, that a sentence of the form *a is F* is true just if what *a* names is in the extension of *F* is not true? Shall we say that the anti-Semite who utters 'I promise to give back the money I took from the \*s' cannot keep his promise (because to keep a promise to *V* requires that one makes true what's said by one *I will V*), or shall we instead tell a novel story about what it is to keep a promise? One worries that whatever we say, we will end up complicating or reformulating the semantic account of our language, replacing relatively simple accounts of reference, truth, promise keeping, and so on with more complicated ones. And this is a theoretical cost.<sup>37</sup>

There are several things to be said in response to this. The first is that it's not clear that the 'theoretical cost' incurred here is all that large. Is it, for example, all *that* great a complication in theory to say that when a term's use has a performative or evaluative dimension, we need to distinguish between the things of which it is true and the things to which linguistic conventions sanction its application? Is it all that great a complication to go on to say that if these two sets do not coincide—and so there is something defective about the relevant

<sup>36</sup> One might object to what I have said in the following way. You allow us to move from 'B uttered "Smith is a \*" and believed what he said' to 'B believes that Smith is a \*'. The inference is all right if the use of 'Smith is a \*' under 'believes' in the conclusion says what it says when uttered by B. But your point here depends on *denying* that the sentence as used by B says what we say with it in indirect discourse.

In response: Forget slurs for the moment and think about proper names. If Thornton says 'Derrida naps', saying that Jacques Derrida naps, I can echo him and ascribe an assertion to him, saying 'Thornton says that Derrida naps'. This is so even if the 'ways of thinking' of Derrida that Thornton and I associate with the name 'Derrida' are *wildly* different; it is so even if Thornton's way of thinking of Derrida contains elements that I do not ascribe to Derrida and vice versa. The inference mentioned in the last paragraph does *not* depend upon ascriber and ascribee expressing *precisely* the same thing with the sentence used to express and ascribe; rather, it requires (putting things intuitively) that the ascribee's utterance on this occasion can be *translated* by the ascriber's utterance. For further discussion see Richard (1990: chs. 2, 3) and Richard (2006).

<sup>37</sup> Stephen Davis and Tim Williamson made remarks that suggest this sort of objection, though they are not responsible for the particular spin I have given it.

sentence uses—compliance with a promise or an order is to be explained in some way other than in terms of truth?

Secondly: while we of course don't want to needlessly complicate our theories, we also don't want them to be so simple-minded that they distort our picture of what we are theorizing about. The conclusion that what is said in slurring speech is neither true nor false comes rather directly from the everyday observation that to think what someone like B thinks, when he thinks that the person in front of him is a \* or thinks that the person is not a \*, is to think wrongly because it involves misplaced contempt and hostility. As I see it, we can avoid this conclusion only if we oversimplify the phenomena semantic theory is primarily about—thought and its expression—by artificially ignoring the performative and emotive dimensions of thought and talk when we theorize.

## 1.8

Those versed in the ways of analytic philosophy of language are perhaps puzzled by my failure to consider 'the most obvious' explanation of our reluctance to ascribe truth or falsity to what is said by slurring. This is an explanation in terms of one or another 'pragmatic' mechanism of the sort Paul Grice famously made us aware a half century ago. My excuse for deferring discussion of this is that my doubts about such explanations turn on what I think is the proper story to tell about slurring. That tale has been told, and it is time to take up the idea that the right account of the phenomena we've been discussing lies in pragmatics.

Grice's ideas presuppose a distinction between what a sentence's use says ('strictly and literally') and what it implies. The rough idea can be brought out by noting, with Grice, that (normally) to utter 'Bill produced sounds which closely matched the notes of "Sunday, Bloody Sunday"' implies, but does not *say*, that Bill sang 'Sunday, Bloody Sunday' poorly. According to Grice, what a sentence's use says is closely tied to the sentence's literal meaning; an important mark of the distinction is that an utterance might be true (false) when what is implied is false (true), but an utterance's truth-value must be the truth-value of what the utterance says. Not only did Grice provide this distinction; he offered an elaborate theory of how a sentence use might come to imply something it did not say.

Grice used the distinction and the accompanying theory to defend various accounts of the meaning of words that interest philosophers—the little words of logic ('or', 'if', 'all', etc.), 'know', 'true', and some others—against objections that the accounts predicted sentences in which the words occur to have truth-values that pretty much no one thought they had. The crux of the defense was to note that what is implied is often—perhaps most of the time—of much greater conversational moment than what is strictly speaking said. Thus, it is to be expected that we will focus on what is implied, not necessarily on what is said,

in evaluating an utterance. An example: we tend to reject a conditional *If A then B* if we see no connection between the A and the B. Some have said that this shows that to assert a conditional is to assert *inter alia* a connection between its antecedent and consequent. This is contrary to what the logician teaches, that *If A then B* comes to no more or less than *Either not A or B*. Grice observed that we suppose (absent contrary evidence) that a speaker has reason to say whatever he says; but one has reason to assert *Either not A or B* only if one has some reason to think that A's being true in some way brings B's being true in its wake. Since it is manifest to speaker and hearer that one would assert the conditional only if one thought there was this sort of connection, asserting the conditional implies that A and B are so connected. Thus, the intuition that uttering *If A then B* conveys the existence of such a connection is just what we would expect *given* the logician's account of its meaning and the just-mentioned facts about conversation.

Thus goes the Cliffs Notes summary of Grice's 'Logic and Conversation'. What does all this have to do with epithets and attitudes? Well, someone might propose that when S is a slur and N is its neutral counterpart, *He is an S* is true just in case *He is an N* is true. But a use of the former sentence has an offensive implication, in Grice's sense of implication, on which we understandably fix. We are so concerned to reject it, the suggestion would be, that we miss the fact that the offensive implication is a by-product of the assertion of something which is true. What to make of this suggestion depends on the details; I will consider the two most likely stories.

The first is that uses of slurs carry what are sometimes called 'particularized conversational implicatures'—ones which are made in a one-off fashion (as opposed to the putative 'generalized' implicature about A's bringing B in its wake carried by the conditional *If A, B*). A standard example of a particularized implication is this: I ask 'Where can I buy gas?'; you say 'There's a gas station a mile south of here'. Given the context in which your utterance occurs, you imply, but do not say, that one can purchase gas a mile south of our location. The implicature occurs (as do all conversational implicatures) because, given the assumption that you are following 'the rules of conversation' (try to say true things you have evidence for; be orderly; be perspicuous; be helpful; etc.), the best explanation of your utterance involves supposing that you want me to infer, from your making it, that one can buy gas a mile south of our location.

Why do I say it is implausible that slurring involves conversational implicature? Well, this sort of implicature occurs, so to speak, because it solves a problem posed by someone's utterance: What must they want us to think, given what they said and that they are 'playing by the conversational rules'? The problem is solved by seeing what follows from the assumption that you are playing by the rules, that you said what you did, and other things obvious to everyone. But whatever interpretive problems slurring might cause, it is not clear how appeal to maxims of conversation will help solve them. For nothing much seems to follow from the facts that you insulted the slur's targets by slurring them and that

you are trying to be conversationally relevant, helpful, orderly, and so on—at least nothing much follows that couldn't be inferred from one or the other of these facts alone. In particular, simply to say *He is an S* (S a slur) is to express the thought that the relevant person is an S, and to convey that one accepts this thought. To think *this* thought is to think the person worthy of contempt. So *simply* to say that a particular person is an S is to represent that person as contemptible. It is what the person says and thinks, when he slurs, that we want to reject. We don't need to drag a Gricean mechanism in to understand this. One of course wants an explanation of how the speaker could speak as he did, when he slurs another. It seems a bit meshuga to look for it in the principles of conversation.

One might somewhat more plausibly say that while it is 'part of the meaning' of a slur that its target is contemptible (or inferior, or . . .), this is not part of the truth-conditional content of such terms; the term of art here is 'conventional implicature'. The idea of conventional, non-truth-conditional implicature is grounded in Grice's remark in 'Logic and Conversation' that

If I say (smugly), *He is an Englishman; he is, therefore, brave*, I have certainly committed myself, by virtue of the meaning of my words, to its being the case that his being brave is a consequence of (follows from) his being an Englishman. But while I have said that he is an Englishman, and said that he is brave, I do not want to say that I have *said* (in the favored sense) that it follows from his being an Englishman that he is brave.<sup>38</sup>

Grice offers no explanation of why he does not want to say this, but one can see why he might: To use 'therefore' in the way in question might be said to *perform* an act of drawing a conclusion; drawing a conclusion is not *saying* that one is doing so, or saying that one is doing so validly. Because the purpose of 'therefore' is to indicate that one is drawing a conclusion, this is 'part of its meaning', and the implicature is thus rightly said to be conventional. Some might think that something similar is true of 'but': its purpose, it might be said, is to *draw* a contrast; drawing a contrast is not saying that one is there, but it does in some important sense imply that there is one to be drawn.

In the case of 'but' and 'therefore', the argument for a conventional implicature is that the words are marked for *doing* certain things beyond the humdrum semantic tasks of referring, predicating, quantifying, and so on. It is thus manifest that their use imparts some information, that one is doing the thing for which the words are marked. But, of course, if it is manifest that one is (sincerely) drawing an inference or contrast, it is manifest that one thinks the inference is valid, the contrast there to be drawn. Thus, the use of these words conveys the (perhaps mis)information that the inference is valid, the contrast genuine. Since the information conveyed is not encoded (simply) by the reference, predication,

<sup>38</sup> Grice (1967: 22).

quantification, etc. which occurs as one utters the sentence, that information is not 'part of what is strictly and literally said'.

One might think that the racial epithets fit this pattern quite well. After all, they are conventional means of expressing derogatory attitudes towards their targets. They are marked for *doing* something beyond the humdrum semantic tasks of referring, predicating, and so on. So, one would expect, their use imparts (non-*semantically*) certain information. Just as drawing an inference, from A to B, conveys that one is justified in doing so—conveys that the inference is valid—so, one might say, displaying contempt for someone on the basis of race or ethnicity implies that one is justified in doing so—implies that they are contemptible because of race or ethnicity. We thus reject the utterance in which the slur is used. Not paying close attention to the distinction between what is said and what is implied, we mistake this for rejecting what the sentence 'strictly and literally says'.

I do not wish to deny that there can be—that, indeed, there is—information conventionally associated with a (use of a) sentence, conveyed by it(s) use, which does not enter into determining whether the sentence(s) use is true or false. I do not want to deny that when someone utters *He is an S*, S a slur, linguistic convention (in particular, the fact that the slur is marked as a device of derogation) puts the hearer in a position to see that the speaker thinks, and would be happy to have the hearer think as well, that the target of the slur is contemptible. But I do deny that it is by appeal to this fact that we best explain our reluctance to say that (when the slur is aimed at its target) what is said is true.

A slur is a device made to denigrate, abuse, intimidate, and show contempt. Such is its conventional potential. But because of this it is also a device that is used to portray, to represent its targets. The racist *thinks of* the targets of a slur S as Ss. (Indeed, he wants us and the targets themselves to think of them(selves) in this way.) There is of course a connection amongst the functions of the slur. To refer to someone as an S (S a slur targeted on T) is to *show* contempt for him on the basis of his being a T; to think of someone as an S is to *think* of him as contemptible. The thought that one is having, when one thinks *He is an S*—the thought that sentence, in virtue of what it means, gives vent to—is a thought in which the relevant individual is represented as contemptible. So, at any rate, I have argued in the second half of this chapter. *That* thought must be rejected—not by asserting its negation (and thereby continuing to represent the individual as an S), but by its outright rejection.

It is our recognition of this that leads us to reject the thought that someone is an S. That is, it is our rejection of the thought that *He is an S*—what the sentence *says*, in as strict a sense of 'says' as you like—that is responsible for our reaction. But a sentence that strictly speaking says something that is not to be accepted as true—what such a sentence strictly speaking says is not true.

## 1.9

I have, in some ways, pronounced the intuition with which we began—that slurring speech cannot be true—both right and wrong. I have, in some ways, said that Gibbard's verdict on such speech—that it may be true, though 'offensively couched'—is both wrong and right. For I have suggested that in different situations we might individuate the way of thinking expressed by a slur differently. If we include the contemptuous attitudes annexed to that way of thinking—something I have been arguing it is possible and permissible, sometimes even required, to do—we have a thought we should not call true or false. But it may be possible, in some situations, to think of this way of thinking as only accidentally involving the relevant attitudes. If we can do this, we can presumably ascribe truth to a thought employing this way of thinking—if, at least, there is nothing wrong with it beyond the attitudes conjoined to it.

I close by observing that this last possibility is very often one that seems not to be open to us. It is just not open to me to unilaterally detach the affect, hatred, and negative connotations tied to most slurs and use them interchangeably with their neutral counterparts. Whether I like it or not, I am unable to use most slurs without showing contempt for their targets. I would say that the same thing is true of approving of or ascribing truth to uses of slurs by others. It might be that in another place and time, if few or no people had the attitudes now associated with a slur, I could ascribe truth to what someone said using it without in effect endorsing the attitudes expressed, just as Blackburn can today call something lustful without thereby condemning it. But at this juncture, at this particular point in history, I can't.<sup>39</sup>

<sup>39</sup> It may seem that I have said nothing which applies to most racist talk. Most racism is achieved without the signaling of contempt with specific devices. The professor who remarks, when a colleague laments the paucity of minority students in the philosophy program, 'Well, one has to have strong quantitative abilities to do well in philosophy, and on the whole they don't' would normally and correctly be taken to have made a racist remark. (The example is a variant of one due to Cara Spencer.) Does anything I've said apply to this kind of case?

I think so. Assuming that 'they', as used here, contributes a certain way of thinking of a group of people to what's said, the question to ask here is whether the way of thinking is like those associated with a slur in the relevant regard: that is, is it a way of thinking of someone which 'fronts' denigration and contempt in the way a slur does? If so, the right thing to do is not logic chop or engage the man on a descriptive level (arguing, say, that he is confusing training in standardized test taking with intellectual ability). The right thing is to reject the way of framing the targets of the thought. Some claims are claims we ought to reject not because they are false but because to seriously take them or their denials up requires thinking of the world awrong.