

THE SEMANTICS OF RACIAL SLURS: USING KAPLAN'S FRAMEWORK TO PROVIDE A THEORY OF THE MEANING OF DEROGATORY EPITHETS

JOSEPH A. HEDGER

jhedger@syr.edu

Department of Philosophy
Syracuse University

ABSTRACT. In this paper I adopt Kaplan's framework for distinguishing between *descriptive* and *expressive* content. Racial slurs are an especially difficult challenge for truth-conditional semantics because of their projection behaviors. That is to say, the offensive content of slurs "scopes out" of logical operators. I argue that racial slurs express contempt and lack descriptive content, so that many sentences containing slurs are not truth apt. My theory accounts for the intuition of the ordinary speaker who refuses to assent to the truth of a sentence containing a slur, but accepts the same statement made using a neutral counterpart of that slur. Weaknesses of rival theories (including those of Williamson, Hom, and Richard) are briefly discussed.

Keywords: slurs, expressivism, David Kaplan, semantics, epithets, projection

1. Preliminaries

Before we get started, I'd like to just remind ourselves that, as we all know, slurs are offensive words that can hurt people, and as we'll discuss later, their offensiveness projects through almost any type of linguistic construction. In fact, not only would I rather not *use* any slur words, but I'd rather not even *mention* any. In the past I've used less offensive slurs as examples, but part of the problem with that tactic is that the more odious examples are actually better at priming the kinds of intuitions that I rely on. So I'll use a capital 'S' as a schematic letter standing in for a racial, ethnic or religious slur. Since we are all unfortunately aware of many slur words, I invite you to think to yourself of an example that you find particularly offensive, and then make your judgments about that sort of example statement. Also, this paper concerns only *slurring uses* of slur words. There are

other interesting phenomena, such as appropriation – which is where a targeted group takes on a slur word to address other insiders of that particular group – but I won't have anything to say about those sorts of uses in this paper.

2. Truth Conditional Semantics

Semantics is the study of the *meaning* of strings of language. There's a tradition in linguistics and philosophy which places a large emphasis on *truth conditions*. The basic picture, stemming from work done by Frege, Carnap, Kripke, Kaplan and others, is that (since language is compositional) the meaning of a statement is determined by the meaning of its parts. Take, for example

(1) David is intelligent.

The extension of 'David' will be a particular individual, and the thought or proposition expressed by (1) predicates intelligence of David. So, (1) is true just in case David is intelligent, and false otherwise. According to Frege and the tradition following him, the extension of (1) will be a truth value. One popular way of understanding the intension of (1) is as a function from possible states of affairs to truth values. Thus, meaning either is, or is what determines, a truth value. Since extensions are truth values, logical operators function the same way in semantics that they do in classical logic. Thus, for example, someone who asserts (2) also asserts (1) but someone who asserts (3) does not assert (1). This is because the truth of (2), but not the truth of (3), depends upon the truth of (1). (4) is of course false whenever (1) is true.

(2) David is intelligent and Canada is cold.

(3) If David is intelligent, then so is Judith.

(4) It is not the case that David is intelligent.

3. Projection Behaviors of Slur Words

However, the projection behaviors of slur words pose a challenge to truth conditional semantics. That is to say, the use of a slur remains offensive even when that slur word is embedded under negation, as the antecedent of a conditional, in a question, etc. Thus, someone who uttered (5) would be uttering something clearly offensive. In this sense (5) contrasts with (3), because asserting (3) doesn't commit the speaker to predicating intelligence of David or Judith. (3) could still be true even if neither David nor Judith is intelligent. However, someone who utters (5) has still nonetheless said something offensive.

(5) If Obama is an S, then so is his wife.

Still, maybe it's not such a bad thing to expand our semantic theory beyond truth conditions. This approach has arguably worked well for a limited case of expressions – yes/no questions, statements, and commands – but it appears poorly equipped to handle the semantics of other types of expressions. I want to suggest that a framework used by David Kaplan (2005) can be extended to the semantics of slur words (even though in the end I will disagree with Kaplan's specific analysis of slurs).

Kaplan borrows a technique from truth conditional semantics in order to introduce his distinction between two different kinds of semantic content. If two statements are identical in semantic content, then we should be able to make a valid inference from one to the other. For instance, take (6) and (7). Since 'polite' and 'courteous' are synonyms and roughly equal in semantic content (let us assume), we can validly infer (7) from (6) or (6) from (7). However, as Kaplan (2005) notes, intuitively we cannot validly infer (9) from (8). On the other hand, validly inferring (8) from (9) seems ok.

(6) David is polite.

(7) David is courteous.

(8) Kaplan was promoted.

(9) That damn Kaplan was promoted.

4. Descriptive Semantic Content and Expressive Semantic Content

Kaplan explains this difference by arguing that (9) contains the semantic content of (8), but also something extra, viz. a display of an attitude of contempt on the part of the speaker. (8) is composed of purely *descriptive* content – it represents the world as being a certain way, and as such can be either true or false. However, expressions such as 'damn' (used as an adjective), 'ouch' or 'oops', contain only *expressive* content. These expressions merely display an attitude of the speaker, and as such are not truth-apt. When someone utters "ouch," they express (roughly) that they have just experienced a sudden pain, but it makes no sense to say that what was said is true or false.

5. Two Categories of Derogatory Epithets

Armed with this distinction, we can now distinguish two different categories of derogatory epithets in general. As Mark Richard notes, although it may be difficult to specify the content in practice, we all know what it means to call someone an 'asshole.'

(10) My neighbor is an asshole.

To say (10) is to say something *about* one's neighbor. (10) describes the speaker's neighbor as well as expresses a contemptuous attitude of the speaker. On the other hand, someone who utters (11) doesn't describe their neighbor at all, but merely expresses an attitude of contempt towards them.

(11) My neighbor is a fucker.

Although we might disagree here about specific examples, I want to argue that there is at least some sub-class of derogatory epithets which lack any descriptive content whatsoever. One reason I have for thinking this is that some years ago I learned some Spanish by speaking with some co-workers in a restaurant and looking words up in a Spanish-English dictionary. When I would look up certain cuss words I noticed that different dictionaries would give very different English expressions as translations. It occurred to me that a vast number of derogatory expressions (particularly those considered most offensive) don't obviously differ in meaning.

I think that this is because these epithets have the same expressive content, but lack descriptive content altogether. Consider, for instance, the person who, while working on a car, hurls a wrench in frustration and yells "Fuck!" It would be odd to claim that this person has described the wrench or the car as being a certain way. It would be equally odd to say that 'fucker' is semantically ambiguous between a use such as this and the use typified in (11).

Another consideration in favor of this reading is that we have no basis for arguing with someone who utters (11). We may feel that the speaker's hostile attitude is not warranted, but we have nothing to say which is capable of demonstrating that the neighbor is not a fucker. This is because 'fucker' doesn't describe the neighbor or say anything about him. On the other hand, if we know the utterer's neighbor, we could reason that the neighbor's pleasant disposition means that the neighbor is *not* an asshole. The fact that we have a basis for arguing against the claim made by (10) but not the claim made by (11) supports the idea that (10) contains descriptive content but (11) does not. Consider also the contribution made by 'blasted' in (12).

(12) The blasted TV is not working.

(13) The TV is not working.¹

Intuitively, this word doesn't contribute anything to the descriptive content of the sentence, and it doesn't obviously have different truth conditions from (13). What I want to urge is that if we treat slurs in a similar fashion, then this is the best way to account for many of our ordinary intuitions concerning statements containing them. The view here offered for your

consideration is that slurs contain merely expressive content – i.e. they display an attitude of contempt on the part of the speaker toward their targets – but they lack an extension, and hence don't make a truth-apt contribution to semantic content.

For example, we would all assent to the truth of (14). However, since we are not racist, we would feel extremely uncomfortable assenting to the truth of (15).

(14) Obama is the first black President of the U.S.

(15) Obama is the first S President of the U.S.

These sentences differ markedly in their meaning. Frege called the meaning of an indicative statement the “thought” of the sentence. Sentence (14) expresses a thought which is true, and can be accompanied by feelings of admiration and pride, for instance. Sentence (15) expresses a racist thought, which we wish to condemn and which we refuse to agree with. On my view, (14) is true but (15) is not even truth apt, since it fails to offer sufficient descriptive content to predicate anything of Obama.

6. Slurs and Satisfaction Conditions

Next, we shall have to say something about why slurs should have satisfaction conditions if they lack an extension. Consider (16):

(16) That is an elephant.

If someone utters (16) while pointing to a giraffe, then that person has obviously said something inappropriate. This inappropriateness can be explained by the fact that a giraffe does not fall under the extension of the word ‘elephant.’ We might have similar intuitions concerning (17) when S is a slur normally used to target blacks, and the utterer of (17) points to, say, a person of Swedish descent.

(17) That person is an S.

However, notice that this utterance still manages to be offensive, and that a Swede who felt that the speaker was expressing contempt toward him would not thereby be making a linguistic error.

7. Honorific Titles

I think that this can be explained by the fact that slurs are the other side of the coin of what we might call honorific titles, such as ‘sir’ or ‘miss.’ Just as slurs function to display contempt, honorifics function to display respect. Intuitively, honorific titles also lack an extension. For example, if we thought

that all males fell under the extension of ‘sir,’ then we would have to say that Adolf Hitler should be addressed as ‘sir.’ If we said that the extension of ‘sir’ is any male worthy of respect, then we would have to say that Hitler’s generals made a linguistic error when they addressed him as “Sir.” Either of these theories would also have the result that an utterance of (18), mistakenly made by a server in addressing a woman, would thus be a false utterance. None of these results is satisfactory.

(18) Your dinner is ready, sir.

I think we should say instead that ‘sir’ is composed of purely expressive content, and functions to display the respect of the speaker. Hence, an officer who fails to address his superior as ‘sir’ thereby demonstrates a lack of respect, but does not commit himself to an affront to the officer’s manhood.² If this is right, then honorifics and slurs are alike in having satisfaction conditions of a sort but lacking descriptive content. This would mean that honorifics and slurs are opposite sides of the same coin, which also feels intuitive. To sum up our conclusion thus far, the truth about racial slurs is that they are not truth apt.

8. Rival Theories

I’d like to close by briefly arguing against some rival positions put forward by other philosophers. The case of slurs is a difficult problem, and each theory has its costs. I’d like to consider a few costs of other views which I deem too high to accept. First, David Kaplan holds that slurs contain both descriptive and expressive content.³ Since for him a statement is *true-plus* just in case the descriptive content is accurate and the attitude expressed is genuinely held by the speaker, then his theory results in the possibility that statements containing racial slurs are not only true, but true-plus. However, this idea doesn’t conform to ordinary intuitions, as we have already noted. Furthermore, Kaplan’s resultant theory (by itself) can’t explain why slurs shouldn’t be used. (Kaplan already accepts the framework used here and the semantics of honorifics which I use as a comparison, so I think that he should just say what I say.)

I think that logic, which is the basis for truth-conditional semantics, is meant to be the science of reasoning, and should adhere to norms of good reasoning. Part of my motivation for wanting a semantic theory which doesn’t allow racist statements to be true or to be validly inferred from neutral statements is because this amounts to our accepting these sorts of statements and inferences as somehow appropriate. I suspect that similar motivations lie behind the theories of Mark Richard and Christopher Hom (more on their views below).

Descriptive semantic content is evaluated in terms of truth. Since truth does not apply to expressive content, we need to find some other dimension by which to evaluate the appropriateness of expressive semantic content. Kaplan considers expressive content apt just in case the attitude expressed is one which is genuinely held by the speaker. Therefore, the use of a slur will be appropriate on Kaplan's theory whenever the speaker genuinely holds an attitude of contempt. Hence, on his view, many statements containing slurs will be semantically appropriate – or, in his technical terminology, “true-plus.” This is unacceptable.

Christopher Hom (2008) holds that a racial slur means *ought to be subject to the set of discriminatory practices P because of having the set of negative properties N, all because of being F*. Here F designates the appropriate racial, ethnic or religious group, P is specified by the racist practices of a given institution of racism and N is the set of negative properties assigned to F by the racist ideology of the given social institution. Since this predicate won't properly apply to anyone, then any statement of the form ‘X is an S’ will be false for any name X and any slur S. However, notice that Hom is in an uncomfortable position when it comes to analyzing statements such as (5). If ‘Obama is an S’ is false, then that makes sentence (5) true, according to classical logic and truth conditional semantics. However, intuitively, uttering (5) is just as offensive as uttering merely its antecedent.

On the face of it, Mark Richard (2008) presents a view very similar to the one offered here. He agrees with what we have said about the truth or falsity of statements containing slurs. On Richard's analysis, slurs represent their targets as *worthy of A because F*, where A is an attitude such as contempt or hatred of the speaker and F is being a member of the racial, ethnic or religious group specified by the slur. However, Richard argues, no one can be worthy of contempt merely because of being a member of such a group. Hence, all statements of the form ‘X is an S’ are neither true nor false, and ought to be rejected in Richard's technical sense of that term. However, a crucial difference between Richard's view and mine is that on my view ‘X is an S’ is not truth apt, while on Richard's view it is a truth apt statement which lacks a truth value, comparable to a liar sentence or a statement with a vague predicate (e.g. ‘Joe is bald’ where Joe's baldness is a borderline case). Hence Richard is committed to truth value gaps, a semantic ambiguity and counter-intuitive reading of standard logical operators, and (ultimately) truth relativity. My view is much simpler and contains a lot less baggage, and does just as well at explaining ordinary intuitions. Everyone has to accept purely expressive language strings such as ‘damn’ (used as an adjective), ‘ouch’ and ‘oops’ anyway. My account also allows us to separate the true, descriptive content from the racist, offensive content in

a statement such as “That S Obama has de-privatized student loans,” which Richard’s theory cannot do.

Lastly, Timothy Williamson (2009) holds that slurs share the extension of their neutral counterparts, but carry a negative conventional implicature about them. Hence ‘Obama is an S’ is true where S just means ‘black’ (for Williamson) but conventionally implicates something derogatory about all blacks. First, we have already discussed why we don’t want a theory which results in such statements being true. Second, note that implicature is a pragmatic feature of speech acts. So even if such pragmatic accounts explain the offensiveness of such utterances, they don’t explain why they express objectionable thoughts or beliefs. Third, if semantics captures *what is said* (as Kaplan and Grice argue), then take the example of a racist professor who says in class “All Jews are S’s.” If I am called upon by a school board to testify as to what the racist professor said, then if Williamson’s account were correct I could accurately report that he said as, “All Jews are Jews.” However, intuitively I haven’t complied with the wishes of the school board if I say this.⁴

Lastly, conventional implicature is a monkey wrench in Grice’s otherwise intuitive and neat distinction between what is said and what is implicated, and is probably therefore best avoided.⁵ However, even if we accept conventional implicature, slurs don’t follow paradigm examples of conventional implicature. To take just one instance, the word ‘even’ is a typical example of conventional implicature. Hence

(19) Even Joe could prove that logical derivation.

says (20) A (and is true/false depending solely upon the truth/falsity of A), but conventionally implicates both B and C:

(20) A. Joe could prove that logical derivation.

B. Other people besides Joe could prove that logical derivation.

C. Of the people under consideration, Joe is the least likely to prove that logical derivation.

According to Karttunen & Peters (1979) and Grice, a felicitous response to a question involving conventional implicature such as (21) assents to the statement semantically expressed but challenges the implicature. However, the same response for Williamson’s analysis of example (22) is not felicitous.

(21) Could even Joe prove that logical derivation?

Yes, but Joe is better at logical derivations than most other people.

(22) Is Obama an S?

Yes, but I don’t think that blacks are contemptible.

Finally, Williamson (2009) and others object to an expressivist account of slurs because they claim that an audience to an utterance including a slur has no basis on which to object if it merely expresses the speaker's own attitude. However, just as an assertion of a belief normally automatically carries an endorsement of that belief as true, we might also say that an expression of an attitude automatically carries with it an endorsement of that attitude. Furthermore, the expression of an inappropriate attitude may also be morally blameworthy. Suppose e.g. that a person in a movie theater viewing a serious documentary about the holocaust was laughing and applauding. Intuitively, that person did something wrong. Interestingly, the thought of such behavior is phenomenologically accompanied by the same queasy sense of unease as one experiences when thinking of a statement containing a slur.

9. Conclusion

Intuitively, "Obama is the first black person to become U.S. President" expresses a *different thought* than does "Obama is the first S to become U.S. President." Furthermore, the ordinary language speaker would not assent to the truth of the second statement, even though the first statement is obviously true. Since a proposition is the part of a statement which is truth evaluable, then the two statements must express different propositions. A correct semantics must therefore treat these two statements as containing different semantic contents, and a correct logic must not allow the second statement to be logically implied nor equivalent to the first. Ultimately, the theories of Williamson, Hom, Kaplan, and even Richard fail because of the insistence that slurs describe or *say something about* their targets. The theory presented in this paper keeps all of our important common sense intuitions intact, offers an explanation for them, and emphasizes the real function of slurs and how they work.

The real value of working on a topic such as this is that once we have come to the end of what traditional truth conditional semantics is capable of handling, we will need to look elsewhere for models which can incorporate more natural language forms. Although it has understandably been of great interest to philosophers, making statements encompasses only a small part of what we do with language. As an initial puzzle case for truth conditional semantics, slurs might be a worthwhile area to begin exploring how we can expand our semantic theories in appropriate ways, in order to get a greater understanding of how language expresses meanings. Indeed, it may be an ideal area to study as an example of meaning that ultimately has little to do with truth. As a model for expressive language, Kaplan (2005)'s framework also looks to be a device worth exploring and expanding. If the sketch I have

provided in this paper looks plausible and interesting, then future work can both sharpen and develop the present theory, and also look to other language forms which may be suitable for a similar treatment. We have already seen some uncontroversial examples of expressions which lack descriptive content, and honorifics may be a class of words which are semantically similar to slurs. Furthermore, expressivist accounts for these categories of expressions don't carry with them the familiar obstacles to other forms of expressivism, such as the Frege-Geach problem.

Admittedly, there are a few minor puzzles resulting from my story which will eventually need to be ironed out. As always, there is more philosophical work to be done. All in all, however, I think that I have provided some reasons to prefer my account to the ones previously presented on the semantics of racial slurs. Not the least of these is that the present theory coheres much better with common sense understandings and intuitions regarding racial slurs. I believe, as Wittgenstein did, that common sense should not be treated like an umbrella, but should be carried into the room as a guiding principle when we philosophize (Diamond 1976, 68).⁶

NOTES

1. This example borrowed from Kent Bach 2006.
2. Thanks to Robin Jeshion for this point.
3. Although not entirely clear from Kaplan (2005), he has confirmed this in personal communication.
4. This points to a general problem for any pragmatic account of slurs: Intuitively, whatever is offensive about a slur is not an aside but part of the "at issue" (i.e. *semantic*) content.
5. This is argued for instance by Kent Bach 1999. Also, I think counterexamples to some of Grice's analyses (e.g. the difference between 'but' and 'and') can be given (see also Kaplan 2005 for a similar argument and putative counterexample which differs from my own, the latter of which is presented in my paper, "The Truth about Racial Slurs").
6. Since I began working on this topic in 2008, I have had the privilege of discussing it with a number of great philosophers. Thanks for helpful conversations and e-mail correspondence with Cheshire Calhoun, Elisabeth Camp, David Curry, John Devlin, Kevan Edwards, Jessica Gelber, Sally Haslanger, John Hawthorne, Mark Heller, Christopher Hom, Robin Jeshion, David Kaplan, Jeffrey King, Bernard Kobes, Gerald Marsh, Thomas McKay, Shyam Nair, Hille Paakkunainen, Ángel Pinillos, Mark Richard, Dawn Starr, Steve Yablo and Timothy Williamson. Portions of this paper were presented at Arizona State University (2009), Syracuse University (2011), SUNY Potsdam (2011), and the 2012 APA Central Division Meeting in Chicago. Thanks are also due to audience members who asked questions and made suggestions.

REFERENCES

Bach, K. (1999), "The Myth of Conventional Implicature," *Linguistics and Philosophy* 22: 327–66.

Bach, K. (2006), "Review of C. Potts, *The Logic of Conventional Implicature*," *Journal of Linguistics* 42: 490–5.

Diamond, C. (1976), *Wittgenstein's Lectures on the Foundations of Mathematics, Cambridge 1939*. Chicago, IL: University of Chicago Press.

Hom, C. (2008), "The Semantics of Racial Epithets," *The Journal of Philosophy* 105: 416–40.

Kaplan, D. (2005), "The Meaning of 'Ouch' and 'Oops': Explorations in the Theory of Meaning as Use," Unpublished ms, presented on numerous occasions.

Karttunen, L. and S. Peters (1979), "Conventional Implicature," in C. Oh and D. A. Dinneen (eds.), *Syntax and Semantics, Vol. 11: Presupposition*. New York: Academic Press, 1–56.

Richard, M. (2008), *When Truth Gives Out*. Oxford: Oxford University Press.

Williamson, T. (2009), "Reference, Inference and the Semantics of Pejoratives," in J. Almog and P. Leonardi (eds.), *The Philosophy of David Kaplan*. Oxford: Oxford University Press, 137–58.

© Joseph A. Hedger