Inflammatory Language

ERNIE LEPORE

Abstract

This is a paper is about a particular subclass of pejoratives, namely, slurs. These are epithets that denigrate a group on the basis of membership alone, e.g., on the basis of race, ethnicity, origin, religion, gender, or ideology. They carrry a characteristic sting, prone to cause outrage and even injury. As to the source of their characteristic sting, the predominant position invokes some aspect of meaning. Some of the few who reject this assumption locate the source of the sting in the taboo status of pejoratives. Others think slurs can sting because of negative associations they carry across time. We challenge both approaches and defend an alternative, for which negative associations are triggered *not* by every token of a pejorative, but rather by certain of its *articulations*.

This paper is sort of a precis of my forthcoming book with Una Stojnic, *On Inflammatory Language: The Linguistics and Philosophy of Pejoratives* (OUP, 2024).[1] Various versions of this paper have been presented in various places, each in the service of either introducing novel data about the distribution of slur terms or trying to account for what makes their usage offensive. The key question guiding these discussions has always been: why are slur terms offensive? The obvious answer is that it is because of a combination of what they mean (or convey). And how this information is encoded (or calculated).

The simplest such proposal is that slur terms are offensive because of what they predicate of their target group (Hom, 2008, 2012; Hom and May, 2013, 2018; Neufeld, 2019). In order for an assertion of (1) to be true, whatever its predicate means should be true of Hermione; in this case, having the property of being a muggle-born wizard.

1.   Hermione is a muggle-born wizard.

One problem with this kind of meaning proposal is that it doesn't permit a slur term's meaning to project. The negation of a predicative sentence is true just in case the predicate does not apply to its subject.

---

[1]   The book by Stojnic and myself is more formal and thorough than this paper could be. This paper is intended to introduce various proposals of what makes slur terms so offensive. Any mistakes or other infelicities are due to me; everything else is based on our joint effort.

63

So, on the current account, whatever is expressed by the predicate in (1)'s negation (2) does not apply to Hermione.

2.   Hermione is not a muggle-born wizard.

But, in this regard, the offense that slurs carry does not behave like a predicative content. Instead, whichever offense (3) carries is carried by (4) as well:

3.   Hermione is a mudblood.[2]
4.   Hermione is not a mudblood.

In lingo, it is said that the offense potential of a slur term takes wide scope over negation.

   At this critical juncture, some theorists conclude that, regardless of whatever the specific meaning of a slur term is, it behaves more like a presupposition than a predicate (Schlenker, 2007; Cepollaro *et al.*, 2019; Cepollaro and Stojanovic, 2016). So, for example, both (5) and (6) presuppose (7); neither can be true if (7) is not true.

5.   John quit smoking.
6.   John has not quit smoking.
7.   John used to smoke.

Although a presupposition account better explains projectability than predicative accounts do, it does not go far enough. This is because there are environments that take a wider scope than presuppositional content. For example, placing (5) in the consequent of conditional (8) can block presuppositional content (7) from taking wider scope than the conditional. Ditto for (9); placing (5) inside an indirect speech report can block presuppositional content (7) from taking a wider scope. Put somewhat differently, the conditional in (8) takes a wider scope than presupposition (7); and the indirect report 'said that' takes wider scope than (7) in (9). This explains why (8) and (9) can be true even if (7) is false,

8.   If John used to smoke, then he quit smoking.
9.   Mary said that John quit smoking, but he never smoked.

In short, if slur terms have meaning, it is not encoded presuppositionally; (10) and (11) behave like (8) and (9): placing a meaning of

---

[2]   I assume sufficient familiarity with slur terms and their effects that these effects can be evoked even without my mentioning them. Instead, I will present relevant data using the fictional slur term 'mudblood', targeting wizards born to non-wizard parents – in the fictional world of Harry Potter, as a placeholder for actual slur terms.

64

a slur term in the antecedent of a conditional or inside an indirect speech or attitudinal report does not block its offense potential.

10.   If muggleborns are inferior on account of being muggleborns, then Hermione is a mudblood.
11.   Draco said that Hermione is a mudblood, but I don't think muggleborn wizards are despicable on account of being muggleborn.[3]

In short, the slur term in (10) and (11) is offensive. Any appeal to meaning to account for this offense potential of a slur term requires an encoding broader in scope than what is needed to account for presupposition distribution. One popular emendation is to appeal to Conventional Implicature items, since they are alleged always to take widest scope (Williamson, 2009; Potts, 2005; McCready, 2010).

The main idea is that, though sentences (12) and (13) agree in truth conditions, the meaning of (13) requires a contrast between height and speed, and so, (14) is incoherent as a matter of meaning since its second clause denies what 'but' means. It's like saying that John is a bachelor who is married.

12.   John is tall and fast.
13.   John is tall but fast.
14.   #John is tall but fast, and there's no specific contrast between height and speed.[4]

In defense of the wide scope claim on behalf of Conventional Implicature items perhaps a better datum is provided by a denial of the first conjunct in (15):

15.   #Joe, who is a spy, is in hiding; and he's not a spy.

No matter where we embed this appositive, e.g., inside a negation, conditional, indirect speech report, modal or any well-formed combination of these, the speaker remains committed to the appositive, namely, that Joe is a spy. So, we ask whether slur terms are, or

---

[3]   Notice that in both (10) and (11) I posit distinct meanings for 'mudblood': in (10), I posit the proposition that muggleborns are inferior on account of being muggleborns as its meaning, and in (11), I posit the proposition that muggleborn wizards are despicable on account of being muggleborn as its meaning. Not only do I not defend either of these proposals, I believe both are false. This raises the sticky question as to whether it is ever possible to articulate its meaning. We will return to specificity below.
[4]   '#' indicates something is off with the sentence grammatically or semantically.

behave like, Conventional Implicature items with respect to taking the widest possible scope.

One problem with this suggestion is that sentences with slur terms do not feel oddly off when their alleged definition is violated. So, although (14)–(16) are all off (i.e., the meanings of their components are incompatible), (17)–(19) remain perfectly acceptable.

16. #[angrily] That goddamned dog is barking again! [cheerily] I'm pleased with the dog.
17. Hermione is a mudblood, but muggleborns are not inferior on account of being muggleborn; they have nothing but my respect!
18. I love mudbloods!
19. (A recovering bigot:) Our entire way of thinking about mudbloods is wrong!

(17)–(19) are all acceptable and each is offensive. It seems obvious that we need to find an account that underwrites a very robust projection of offense; but how? Once, again, we confront the Problem of Specifying Content: it is hard to see what meaning could be the linguistically encoded content of the slur term. Clearly, given the acceptability of sentences (17)–(19), it's hard to imagine what content candidates for 'mudblood' have not been ruled out; indeed, it's hard to rule in any reprehensible state of mind as supplying the specific meaning of a slur term.

The projectability constraint we need is actually much harder to respect than any we have been exploring. Consider sentences (20)–(23).

20. '"John is French" means John is French' does not predicate anything of John.
21. '"John is tall but handsome" means John is tall but handsome' draws no contrast between height and handsomeness.
22   '"Joe stopped smoking" means Joe stopped smoking' doesn't presuppose Joe smoked.
23. '"ouch" means ouch' does not express a state of mind (e.g., pain).

Each of these sentences mentions a true meaning attribution; and although in each attribution the expression to which meaning is being attributed is mentioned, it is also used on the right-hand side in this sense: consider a sentence where the term in question is replaced by a synonym. If the term were only being mentioned the truth value might change, but it doesn't. For example, consider '"bachelor" has eight letters' vs '"unmarried male" has eight

letters' – assuming 'bachelor' and 'unmarried male' are synonyms. But now compare (20)–(23) with (24)–(26)

24.   'mudblood' means mudblood.
25.   'mudblood' has 8 letters.
26.   'mudblood' is a slur term.

Each of these is potentially offensive, and so, were meaning theorists correct, it would follow that the meaning of a slur term projects out of even quotation – a sort of Hyper-projectivity. Indeed, the Opinion letters section in the *New York Times* explicitly discussed the editorial decision to publish a Sunday Review essay 'How the N Word Became Unsayable' by J. McWhorter that spells out the N-word; the essay was prefaced with an editorial disclaimer stating, 'This article contains obscenities and racial slurs, fully spelled out. Ezekiel Kweku, the Opinion politics editor, and Kathleen Kingsbury, the Opinion editor, wrote about how and why we came to the decision to publish these words in Friday's edition of the Opinion Today.'[5]

It is sort of a reductio of the meaning thesis that using slur terms in quotations (or meaning attributions) remain offensive; quotations are universally taken to render the meaning of the quoted item semantically inert. And yet it is undeniable that (24)–(26) can be offensive. But if not meaning, then what's left to account for the offense that audiences can suffer through tokens of slur terms?

The Problem of Specificity together with the Hyper-projectivity Constraint drive some contributors to wonder whether confrontations with slur terms aren't offensive simply because slur terms are taboo words; their tokenings are always prohibited – wherever they might occur – even in quotational and meaning attribution contexts. It is the violation of their prohibition that offends witnesses to their tokenings.

Let's call this non-meaning account of the offensiveness of slur terms Prohibitionism. One obvious problem with Prohibitionism is simply that many taboos fail to generate the same pattern of offense as tokenings of slur terms. No one would deem that Tetragrammaton is a slur term for God. The name is prohibited because of the sacredness of what it names. Other religious taboos illustrate a similar point. Also, historically there was a taboo against tokening of the name of the devil for fear of summoning him. But the devil's name didn't thereby become a slur term; its tokening

---

[5] https://www.nytimes.com/2021/05/09/opinion/letters/n-word.html.

was associated with fear of misfortune, not with offense of the sort that a slur term can produce. Other examples illustrate the same point. In the 19[th] century, Ukrainian, Polish, Lithuanian, and Belarusian were banned in the Russian Empire. These prohibitions didn't engender feelings of offense in anyone when violated. Nor did expressions of these languages become slur terms among Russians or anyone else.

Another serious problem with Prohibitionism is that it looks to get the worry backwards. Intuitively, a slur term is taboo because it is offensive, not the other way around. After all, violations of many taboos don't cause offense (Anderson and Lepore, 2013a,b). Since Prohibitionism states it's the violation of a prohibition that generates the offensive potential, it would seem to lack resources to explain the differences between slur terms and other taboo terms, violations of which have rather distinctive effects.

But there is an even more serious problem with Prohibitionism, namely, what I call the Problem of Inheritance. The pejorative potential of a slur term is 'infectious', carrying over to expressions incidentally matching slur terms in articulation. For example, there is the controversy over the incident of tokening an English adverb that orthographically and phonetically resembles, but is etymologically and semantically entirely unrelated to, the N-word; its inherited offensiveness, which has persisted despite the recognized etymological and semantic independence, has been well documented (Kennedy, 2002, pp. 94–95).

Illustration

As O'Hehir writes: 'We pretty much have dumped that word, because it is so easily misunderstood and other words will do, and also because it carries a permanent taint: The only person who would conceivably use it now would be a snickering, anti-p.c. asshole trying to make an obnoxious point. Do we miss it? I submit that we don't' (O'Hehir, 2020).

Or take a more recent case involving a tokening of a Mandarin demonstrative term ('那个'), which acoustically resembles the N-word. The word was tokened in a language class setting, in describing the correlate of the English demonstrative 'that' in Mandarin. The speaker announced which word of which language is about to be tokened.[6]

---

[6] See also the discussion on https://languagelog.ldc.upenn.edu/nll/?p=23691 for other incidents involving the same Mandarin expression.

Where does this leave us? I am not claiming that the speaker who tokens a slur term is blameworthy (or not), that any offense taken is warranted (or not). Nor am I claiming that the tokening was apt or inapt, necessary or gratuitous in achieving a particular point (whether it be rhetorical, pedagogical, artistic, *etc*.). These questions are downstream from the presence or absence of any offense effect. That is, first comes the offensive speech and then its moral assessment. No one should deny that typical tokenings of slur terms can give rise to negative associations rooted in socio-historical, cultural, and psychological factors; these are open-ended, and so, not content-like. However, what I am claiming is that the trigger of these negative associations is *not* the word itself, but rather standard articulations of the word.

Note that on this proposal it becomes quite easy to account for the Hyper-projectivity Constraint. You cannot token a slur term without articulating it. But if I am right that it is standard articulations of slur terms which cause the offensive sting, then regardless of where that articulation occurs so does the potential to offend, even inside quotes or meaning attributions or even homonyms. And as far as meaning goes, as was illustrated above, we are able to deny any particular negative content without fear of linguistic incoherence because we are not appealing to meaning in order to account for an offensive sting. Indeed, for our purposes, slur terms may turn out be synonymous with their neutral counterparts.

The Inheritance Constraint now being endorsed is akin to what we find with offensive gestures, symbols, and imagery. A gesture or a symbol that accidentally shares a shape with an offensive one is prone to trigger the same offensive effect. So, for instance, consider the Nazi *Hakenkreuz* and corresponding ancient Swastika symbols that carry positive connotations in various Eurasian religions, including Hinduism, Jainism, and Buddhism. The Hakenkreuz's violent and offensive symbolism does not only give rise to taboo and prohibitions – often legally codified – against its displays, but its offensive potential easily transfers to closely resembling symbols. This is exactly the same phenomenon we find with the inheritance of the offensive potential of slurs. This is why the unfortunately sounding adverb remains tainted even once all confusion has been cleared.

I need to emphasize that the effect triggered by an articulation of a slur term can be triggered even if no slur (or word) is tokened; and conversely, a slur term can lose its potent sting when its articulation is sufficiently far off from a standard articulation; or even when the articulation itself becomes less and less used. That is, the status of a slur term can change over time, without a change in meaning. And,

of course, the presence and severity of the offensive potential can vary itself with articulations. Further, I must caution not to confuse tokenings of articulations with tokenings of words (Stojnic, 2022).

I want to end this discussion with some more data in support of the articulation account. I begin with the vivid example of the phenomenon of graphic slurs in logographic languages like Mandarin, where it is possible that only certain written, but not spoken, articulations of a particular term can be offensive. For example, the Mandarin exonym 'Yáo', the pronunciation of which triggers no offensive effect, when written, can receive either an offensive or a non-offensive rendition, depending on the choice of alternative (phonetically indistinguishable) phono-semantic compounds (either featuring the person-radical or the beast-radical). That can be interpreted as a slur term when written with one, but not another, phonetically indistinguishable choice of a character, depending on the semantic component of the phono-semantic compound (see Matisoff, 1986).

It is thus only one of several possible standard articulations, a particular spelling, that triggers the offensive potential. If the offensive potential were tied to either the word or its meaning, this would be rather puzzling – a word retains its identity regardless of the choice of whatever acceptable spelling it is rendered in; and it retains whatever meaning it has regardless of how it is articulated. Our account, by contrast, readily explains this: it is a particular articulation, not the word itself, that harbors the offensive potential.

Then there is the so-called censoring asterisk. Consider (27)–(28):

27.   'You mudbl∗∗d,' Draco yelled.
28.   #'You the M-word,' Draco yelled.

You might think they are equivalent, but they are not. (28) consists of the subject 'you' concatenated with a canonical description for a slur term. But (27) seems to be saying Draco called you something. So, what did he call you? If you are of the opinion that replacing the two 'o'-s with two asterisks derives a new word, then (27) is false. Its only reasonable reading that allows it to be true is to permit 'mudbl**d' to be a non-standard articulation of a slur term that has as its standard articulation 'mudblood'. These kinds of data show how non-standard articulations of bad words need not be offensive.

Conclusion

There is obviously much more to discuss, including the prospects of success of other sorts of meaning theories for slur terms (expressivism

(Potts, 2007; Jeshion, 2013), perspectivalism (Camp, 2013, 2018), and even pragmatic accounts of content (Nunberg, 2018; Jorgensen Bolinger, 2020)). Another key topic without enough space to discuss is *reclamation* of slur terms. The forthcoming book mentioned in the first footnote by me and Una Stojnic takes on these central topics and much more.

The major focus of this discussion has been that it is not slur terms, but their standard articulations that carry offensive potential. This means that, however this potential is determined, it has little to do with semantics or pragmatics, or indeed even with language at all.

References

Luvell Anderson and Ernie Lepore, 'Slurring Words', *Noûs*, 47:3 (2013a), 25–48.

Luvell Anderson and Ernie Lepore, 'What Did You Call Me: Slurs as Prohibited Words', *Analytic Philosophy*, 54 (2013b), 350–63.

Elisabeth Camp, 'Slurring Perspectives', *Analytic Philosophy*, 54:3 (2013), 330–49.

Elisabeth Camp, 'A Dual Act Analysis of Slurs', in David Sosa (ed.), *Bad Words: Philosophical Perspectives on Slurs* (Oxford: Oxford University Press, 2018), 29–59.

Bianca Cepollaro and Isidora Stojanovic, 'Hybrid Evaluatives', *Grazer Philosophische Studien*, 93 (2016), 458–88.

Bianca Cepollaro, Simone Sulpizio, and Claudia Bianchi, 'How Bad Is It to Report a Slur? An Empirical Investigation', *Journal of Pragmatics*, 146 (2019), 32–42.

Christopher Hom, 'The Semantics of Racial Epithets', *The Journal of Philosophy*, 105:8 (2008), 416–40.

Christopher Hom, 'A Puzzle About Pejoratives', *Philosophical Studies*, 159 (2012), 383–405.

Christopher Hom and Robert May, 'Moral and Semantic Innocence', *Analytic Philosophy*, 54:3 (2013), 293–313.

Christopher Hom and Robert May, 'Reference, Inference, and the Semantics of Pejoratives', in David Sosa (ed.), *Bad Words: Philosophical Perspectives on Slurs* (Oxford: Oxford University, 2018), 108–31.

Robin Jeshion, 'Expressivism and the Offensiveness of Slurs', *Philosophical Perspectives*, 27 (2013), 231–59.

Renée Jorgensen Bolinger, 'The Pragmatics of Slurs', *Noûs*, 51:3 (2020), 439–62.

Ernie Lepore

R. Kennedy, *Nigger: The Strange Career of a Troublesome Word* (New York: Pantheon, 2002).

J. Matisoff, 'The Languages and Dialects of Tibeto-Burman: An Alphabetic/Genetic Listing, with Some Prefatory Remarks on Ethnonymic and Glossonymic Complications', in J. McCoy and T. Light (eds), *Contribution to Sino-Tibetan Studies* (Leiden: Brill, 1986), 3–57.

Elin McCready, 'Varieties of Conventional Implicature', *Semantics and Pragmatics*, 3 (2010), 1–57.

Eleonore Neufeld, 'An Essentialist Theory of the Meaning of Slurs', *Philosophers' Imprint*, 19 (2019), 1–29.

Geoffrey Nunberg, 'The Social Life of Slurs', in Daniel Fogal, Daniel Harris, and Matthew Moss (eds), *New Work on Speech Acts* (Oxford: Oxford University Press, 2018), 237–95.

Andrew O'Hehir, 'So Much for Youth Apathy: Student Radicalism Escapes the '60s At Last', *Salon*, 2020. Available at: https://www.salon.com/2015/11/17/so_much_foryouth_apathy_student_radicalism_escapes_the_60s_at_last/.

Christopher Potts, *The Logic of Conventional Implicatures* (Oxford: Oxford University Press, 2005).

Christopher Potts, 'The Expressive Dimension', *Theoretical Linguistics*, 33:2 (2007), 165–98.

Philippe Schlenker, 'Expressive Presuppositions', *Theoretical Linguistics*, 33:2 (2007), 237–45.

Una Stojnic, 'Just Words: Intentions, Tolerance, and Lexical Selection', *Philosophy and Phenomenological Research*, 105:1 (2022), 3–17.

Timothy Williamson, 'Reference, Inference, and the Semantics of Pejoratives', in Joseph Almong and Paolo Leonardi (eds), *The Philosophy of David Kaplan* (Oxford: Oxford University Press, 2009), 137–58.