

Croatian Journal of Philosophy
Vol. XXV, No. 75, 2025
<https://doi.org/10.52685/cjp.25.75.1>
Received: November 1, 2025
Accepted: November 14, 2025

The Evolving Understanding of Slurs: An Inquiry into Meaning and Effect of Slurs

UNA STOJNIC
Princeton University, Princeton, USA

ERNIE LEPORE
Rutgers University, New Brunswick, USA

In this précis, we summarize the key themes and arguments of the Inflammatory Language: Its Linguistics and Philosophy.

Keywords: Slurs; pejoratives; content; taboo; pejorative tone; articulations.

Slurs are powerful linguistic weapons. These expressions denigrate those they target purely on the basis of group membership (e.g., on the basis of race, ethnicity, origin, religion, gender, sexual orientation, or ideology). They are highly inflammatory; slurring someone involves a transgression more serious than a mere insult, so much so that even mere tokenings of slurs often have a full-on taboo status, and are subject to media censorship, and even legislation. They have a “viscerally palpable effect,” a characteristic pejorative sting that makes them prone to offend and harm but also imbue them with rhetorical potency that mere neutral descriptors for the group they target lack. This pejorative effect underwrites the intuition many theorists voice that tokenings of slurs as opposed to the corresponding neutral descriptors might be necessary in certain context in order to achieve a particular goal—political, pedagogical, artistic, and so on.¹ But characteristically, even tokenings that aren’t intended to offend and aim instead to create a rhetorical effect that furthers a political, pedagogical, or artistic goal remain inherently risky. We thus commonly see discussions of whether particular tokenings of slurs in literature, poetry, film, theatre, or in

¹ See, for instance, Camp 2013.

political activist discourse, or within a pedagogical setting, were warranted or apt, or gratuitous or flippant. Indeed, even mere quotations and displays of slurs retain their pejorative sting. Scholarly publications on slurs frequently open with a disclosure that slurs will be mentioned (not used) in the text, followed by an apology for the offense this is prone to cause. Indeed, mere *displays* of slurs are often subject to outright media censorship and even legislation. In *Inflammatory Language*, our question is what makes slurs so inflammatory? What is the source and nature of their characteristic pejorative sting?

In searching the philosophical literature on pejorative language, and in particular, on slurs, prior to this century, there is little to be found. The few extant discussions are mostly about the inferentialism. For example, there are published discussions by Robert Brandom and by Michael Dummett that exploit distributional intuitions most speakers share about slurs (Dummett 1981; Brandom 2002). The reason why philosophers, by and large, avoided theorizing about pejorative language is because everyone—philosopher or not—held it, implicitly or explicitly, that slurs derogate and are prone to offend because of what they mean.

In contrast to the relative neglect of the past centuries, the last few decades have seen an explosion in the philosophical and linguistic discussion of slurs, with a wide range of proposals on offer as to what makes them so inflammatory and gives them the pejorative sting. But this natural pre-theoretical stance, and a *prima facie* plausible theoretical starting point—that the key to understanding the inflammatory nature of slurs lies in their meaning—is retained throughout much of that literature. The rough idea is that slurs have some offensive, pejorative meaning, and so that tokenings of slurs tend to be offensive because they express or otherwise evoke that pejorative meaning. Much of the discussion and disagreement in the literature then lies in characterizing this meaning—whether as descriptive, expressive, or some other kind *sui generis* content—and the ways in which it is conveyed—semantically, as a matter of truth-conditions, or presupposition, or conventional implicature, or pragmatically, as a matter of conversational implicature.

In *Inflammatory Language*, we argue that this unanimity about the meaning approach to the pejorative sting of slurs is radically mistaken. In the book, we critically examine various proposals in the literature that posit pejorative content. Such proposals are numerous—indeed, we are not exaggerating much when we say that for any of the familiar ways of conveying or expressing content—be it as *at issue* truth-conditional content, or presupposed or conventionally implicated content, or perhaps pragmatically implicated content—there is a proposal out there in the literature that argues that that is the way in which the pejorative content of slurs is conveyed. Similarly, there is a whole array of proposals as to what the *nature* of this pejorative content is: some argue it is descriptive, representational content, others that it

is purely affective or expressive, and yet others that it is *sui generis* type of perspectival content that combines both affective and descriptive components.

But regardless of the near consensus that the pejorative sting of slurs is a matter of meaning, and the many content proposals on offer as to the nature of that meaning, we argue that no content account can ultimately be successful. We raise various challenges for different accounts specifically, but we also argue there is a host of challenges that in principle tell against *any* content account one might offer: in short, there is no special pejorative meaning either semantically encoded in, or pragmatically conveyed by, slur terms. Here we briefly summarize just a few of the main challenges we raise in the book.

A peculiar thing that one immediately encounters when inquiring into the meaning of slurs is a lack of non-schematic proposals as to what particular slurs actually mean. This is already reflected in the practice most contemporary dictionaries adopt when listing the meaning of slurs. The O.E.D., for example, offers pretty much the same definition for each slur term: it states that it denotes individuals that belong to a certain group, and that the term is offensive or derogatory (with occasional remarks about the term's etymology, frequency or common usage). This schematic approach is present in most contemporary accounts of slurs that posit special "pejorative" content—either pragmatically implied or semantically encoded. So, to mention just a few concrete proposals, we get accounts that tell us, e.g., that a slur *S* targeting a group *G* predicates of an individual *x*, that *x* ought to be a target of negative moral evaluation because of being a member of *G* (Hom 2008; Hom and May 2013; see also Neufeld (2019) for a related proposal); or that it presupposes that (the speaker believes that) the members of *G* are despicable on account of being members of *G* (Schlenker 2008; Stojanovic and Cepollaro 2015; Cepollaro 2015); or that it conventionally implicates that the speaker has (occurrent or standing) negative affect towards the members of *G* (on account of being members of *G*) or having negative traits stereotypically associated with *G* on account of being members of *G* (McCready 2013; Williamson 2009; Potts 2007). In each case we are given a meaning schema, the instances of which lay out the meanings of particular slurs.

The schematic accounts of this sort—much like the dictionary practice—essentially posit the same meaning for all slurs terms, modulo the group membership. This, however, is intuitively inadequate. There is a great deal of variation in the severity and character of the offensive effect of different slurs—including different slurs that target the same group of people. (Consider, for instance, distinct slurs targeting women, or gay men, or African Americans.) Yet the schematic accounts above would tell us that any two slurs targeting, say, women, have the same meaning: e.g., that they predicate that the subject ought to be the target of negative evaluation because of being a woman, or that they presuppose that (the speaker believes that) women are despicable

on account of being women, or that they conventionally implicate a negative affect towards women, and so forth. This seems woefully inadequate.

This lack of specificity, we argue, is not a coincidence. One might be tempted to overcome it by positing different specific contents along the lines of those specified by the schema for different slurs. So, one might be tempted, say, to argue that different slurs targeting the same group encode or express different negative attitudes or different affects, and that those different attitudes or affects are of different severity and strength. But this reaction is on the wrong track. For one, it's hard to see how any such semantically encoded content could be what underlies the intuitions about the severity and differences in the pejorative sting of different slurs targeting one group; for while it's undeniable that there are such variations in the strength and severity, competent speakers routinely disagree over which slurs for a particular group are "worse" or more severe, and why.² But these disagreements do not display linguistic ignorance or incompetence with *meaning*. And there can be changes in the nature and severity of a slur's offensive sting over time—as well as differences across dialects; but such changes don't suggest a change in *meaning*.

Moreover, for any bit of specific semantic content one might posit, one can consistently deny that content while predicating a slur. This is already evident even with the bare-bones schematic proposals above; viz.:³

1.
 - a. Hermione is a mudblood, but she ought not be subject to negative moral evaluation on account of being muggle-born.
 - b. It's just false that mudbloods are inferior/depicable on account of being muggleborn. I have nothing but respect for them!

² Consider, for instance, this discussion over the comparative severity of distinct slurs for woman engaging in casual sex: <https://www.thestudentroom.co.uk/showthread.php?t=2171660>. One can easily find an abundance of further examples of debates of this kind on the web.

³ 'Mudblood' is an extremely offensive slur in the fictional world of Harry Potter, targeting wizards born to non-wizard parents. 'Muggle-born' is its neutral counterpart. We stick to presenting examples involving fictionalized slurs, rather than quoting actual slurs; this is not because we advocate silentism or blanket censorship of such tokenings. We emphatically do not. But as we'll see, even mere quotations and displays of slurs (or more precisely, their articulations), carry the pejorative sting. That doesn't mean all such tokenings constitute transgressions: we argue the sting can be not only weaponized to harm, but also used to achieve various potentially worthy rhetorical, artistic, pedagogical, political or other goals. However, we prefer not to judge whether any such tokening on our part would be apt or warranted. We trust our readers can test their intuitions against non-fictionalized counterparts of our examples featuring slurs they are familiar with. For a more thorough discussion of these issues, see Ch. 1 of *The Inflammatory Language*.

- c. Hermione is a mudblood, and I *love* mudbloods! They are my favorite wizards!
- d. Mudbloods have no negative features due to being muggle-born.

By contrast, denying the predicated, presuppositional, or conventionally implicated content, be it descriptive or expressive, *does* lead to linguistic incoherence:

2.
 - a. # The period between October 18th and November 1st is two weeks, but not a fortnight.
 - b. # John stopped smoking, but he never used to smoke.
 - c. # John is tall but fast, and there's no contrast between height and speed.
 - d. # Joe, who is a spy, is in hiding; and he's not a spy.
 - e. # [Angrily:] That goddamned dog is barking again! [Cheerily:] I'm pleased with the dog.

Making the schematic content more specific does nothing to alleviate this problem: take whatever description, attitude, or affect you think a slur encodes—we submit you can consistently predicate the slur while denying that description, or affect, or attitude.

One might at this point be tempted by an account according to which the pejorative sting of a slur is a matter of conversationally implicated or otherwise pragmatically implied content, for such content is cancellable, and thus deniable.⁴ But this temptation should likewise be resisted. First, even a proponent of a pragmatic account has to give a plausible account of what the specific content implicated or otherwise pragmatically conveyed by tokenings of slurs is. Such account must respect the variation in the strength and severity of the offensive potential of different slurs (including different slurs for the same group), and at the same time posit content that is plausibly grasped and tracked by ordinary speakers. We are pessimistic about the prospects of meeting this challenge. But more importantly, while predications of slurs are consistent with denials of any particular bit of descriptive or affective content, as the data above show, the pejorative sting of slurs is precisely *not* cancellable. Indeed, that is already plain to see in examples (2a)–(2d). While the examples illustrate that one can consistently predicate a slur while denying the supposed pejorative content, at the same time, the tokenings of slurs in (2a)–(2d) still retain their offensive sting: (2a)–(2d) are still prone to cause offense, and have radically different effect than their counterparts with the slur replaced with the corresponding neutral descriptor. (We invite our readers to test their judgements with non-fictional slurs they are familiar with.) In short: any bit of (descriptive or affective) content is cancellable, but the pejorative sting is not—hence, the pejorative sting is not a matter of content.

⁴ See, e.g., Nunberg 2018; Jorgensen Bolinger 2017.

Indeed, one of the key challenges for content accounts is what we call the “hyperprojectivity” of the pejorative effect: not only does it remain even under negation (3a), but it persists in—i.e., projects out of—a wide-variety of embeddings, including presupposition filters and plugs ((3b)–(3c)):

3.
 - a. Hermione is not a mudblood.
 - b. If muggleborns are inferior on account of being muggleborns, then Hermione is a mudblood.
 - c. Draco said that Hermione is a mudblood, but I don’t think muggleborn wizards are despicable on account of being muggleborn.

By contrast, negated predications don’t predicate anything of the subject (or anyone else), while presupposed content doesn’t project out of filters when the presupposition is locally satisfied, nor does it project under plugs. For example, (4a) doesn’t predicate being Italian of Mary, while (4b) and (4c) don’t presuppose that John used to smoke:

4.
 - a. Mary is not Italian.
 - b. If John used to smoke, then he stopped smoking.
 - c. Mary said that John stopped smoking, but he never smoked.

Indeed, the pejorative sting notoriously projects even out of environments that normally render any type of meaning inert—the purely quotative uses, and indeed even mere displays of slurs. In purely quotative environments, an expression is merely mentioned, not used to express its meaning; so, for instance, semantic content is rendered inert in meaning attributions (both on the right-, and left-hand side) (Anderson and Lepore 2013):⁵

5.
 - a. “David is Jewish’ means David is Jewish” predicates nothing of David.
 - b. “John is tall but handsome” means John is tall but handsome draws no contrast between height and handsomeness.
 - c. “John, who is my friend, is happy” means that John, who is my friend, is happy” doesn’t convey anyone has friends.
 - d. “Joe stopped smoking’ means Joe stopped smoking” doesn’t presuppose Joe once smoked.
 - e. “ouch’ means ouch” does not express a state of mind (e.g., pain).
 - f. “vous’ means vous” does not commit the speaker to addressing anyone in a formal/ polite manner nor does it signal politeness.

⁵ Examples (5a)–(5d) are from Anderson and Lepore (2013); (5e)–(5f) are our addition.

Moreover, mere quotation generally renders meaning inert:

6.
 - a. ‘Italian’ has 7 letters.
 - b. ‘Italian’ is an English word.
 - c. ‘Italian’

By contrast, the pejorative sting of a slur persists even in such environments (as suggested by well-documented real-world incidents, empirical studies, as well as editorial practices):⁶

7.
 - a. ‘mudblood’ means mudblood.
 - b. ‘mudblood’ has 8 letters.
 - c. ‘mudblood’ is a slur term.
 - d. ‘mudblood’

On the whole, then, we are pessimistic about the prospects of pejorative content accounts. The data above are damning for any such account.

While content-based accounts are by far the most prevalent in the literature, there is an alternative minority position: explain the pejorative sting of a slur not as a matter of its meaning, but through some other feature of the word. So, for instance, Prohibitionists argue that slurs are prohibited, taboo words, and it is the violations of the prohibition that their tokenings constitute that generate their pejorative sting (Anderson and Lepore 2013). The Pejorative Tone Account, in turn, maintains that slurs carry “pejorative tone,” understood in a broadly Fregean sense: they give rise to a wide-range of open-ended, pernicious associations due to a variety of socio-historical, cultural, psychological, and other factors, and it is because of these associations that slurs have pejorative sting and taboo status (Lepore and Stone 2018; c.f. Frege 1897/1979).

While these accounts might fare better insofar as they don’t attempt to explain the pejorative sting in terms of meaning and thus can offer an explanation as to why even mere quotes and displays of slurs retain the pejorative sting, they still fall short. In the book, we raise a range of problems and challenges for both Prohibitionism and the Pejorative Tone Account. In the interest of space, here we shall only point to the most pressing issues that no account that ties the explanation of the pejorative sting of slurs either to their meaning or the words themselves can adequately address.

To see the main issue, we begin with the phenomenon we’ve dubbed “inheritance”: the pejorative sting of slurs is not only hyperprojective, but it is infectious in that it carries over even to expressions accidentally matching slurs in articulation: in their orthographic and phonological form. The phenomenon can be illustrated by the debate and contro-

⁶ See, e.g., Fasoli, et al. 2015; Carnaghi & Mass 2008; Cepollaro, Sulpizio, Bianchi 2019. For a more in-depth discussion of the relevant empirical studies, real-world cases, and editorial practices, see Ch. 6 of *The Inflammatory Language*.

versy over the offensiveness of an English adverb which bears phonetic (and orthographic) similarity to the N-word, but is otherwise etymologically, and semantically unrelated (Kennedy 2002: 94–95). Kennedy reports fiercely divisive opinions surrounding the infamous incident, where a white speaker, addressing a largely African-American audience chose this (rarely used) adverb instead of any of the many (much more frequently used) synonyms (e.g., ‘ungenerous’, ‘stingy’), sparking outrage. Even if one thought that the original outrage was somehow a result of an etymological or semantic confusion, the controversy continues alongside a widespread recognition that the similarity is a matter of mere orthographic accident (to borrow Quine’s apt phrase (1953: 67)), and the adverb remains tainted; as O’Hehir puts it, the adverb now “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” (O’Hehir 2020).

Indeed, the phenomenon of inheritance arises even in cases where there is *no* temptation to post an etymological confusion. Just one illustration is a relatively recent case where outrage and offense were sparked in a class setting by an instructor mentioning (not using!) a Mandarin demonstrative term (‘那个’), which acoustically resembles the N-word. In this case, even at the time of the incident, there wasn’t any confusion about which word was being tokened. The tokening occurred in a context of discussing filler-words, as a part of an explanation that, in Mandarin, this word—which the speaker explained, prior to tokening it, to be a correlate of the English demonstrative ‘that’—is often used as a filler-word. The speaker preceded the tokening with an explicit explanation of which word of which language will be tokened.

Note that we are not interested in the question of whether any of these tokenings constituted a wrongdoing, or whether their speakers were blameworthy; nor are we adjudicating whether any offense they caused was sincere, appropriate or warranted. But it would hardly be necessary to debate these issues, issue apologies, or write an op-ed, let alone numerous ones, on whether these reactions were warranted or appropriate, if it were *obvious* that there was no pejorative sting triggered in these instances.

And it’s worth emphasizing that the phenomenon isn’t isolated or exclusive to a single slur. Indeed, it is quite common. To illustrate with just one more example, consider, for instance, the incident involving a Krispy Kreme doughnuts advertisement that aired in Australia, where the word ‘congratulations’ was rendered with two doughnuts replacing “o” so that a part of the word accidentally matched a highly charged racial slur.⁷ This sparked outrage, the add was removed, and the company issued an apology. Note, here, the orthographic accident is quite

⁷ <https://www.mirror.co.uk/news/world-news/krispy-kreme-apologises-after-racist-30761080>

apparent: no slur was tokened, and the orthographic string (stylized by way of doughnuts) that accidentally matched the articulation of a slur was a proper part of the articulation of the word actually tokened.

Indeed, the phenomenon is a general one that we commonly see with offensive gestures, symbols, and imagery. Consider, for instance, the Nazi Hakenkreuz: its violent and offensive symbolism not only gives rise to taboo and prohibitions—often legally codified—against its displays but its pejorative sting easily transfers to certain closely resembling symbols that carry positive connotations in various Eurasian religions, including Hinduism, Jainism, and Buddhism. This is apparent both in debates over the possibility of “reclaiming” these symbols and “restoring” the original positive connotations, but also in the common attitudes toward them in the West.⁸

We illustrate with an incident involving toy pandas sporting the manji symbol—a good luck symbol that is a mirror image of the Swastika—that ended up in Christmas crackers in Canada, causing widespread outrage, and a subsequent apology and recall from the manufacturer.⁹ This is exactly the same phenomenon, we believe, we find with the inheritance of the pejorative sting of slurs in articulatorily resembling expressions.

The phenomenon of inheritance is difficult to explain on the content-based accounts and the word-based accounts alike: for in tokening the now infamous adverb, or the Mandarin demonstrative, or ‘congratulations,’ no slurs were tokened: and so no slur was even mentioned, let alone used, and thus its (putative) pejorative meaning expressed, or a slur-related taboo violated, or tone triggered.

We instead argue that the pejorative sting of slurs isn’t a matter of their meaning, nor even slur *words* themselves. Rather, we argue that it is (certain of the) articulations of slur words that harbor the pejorative sting. The sting, in turn, is constituted by the open-ended cluster of negative associations rooted in various socio-historical, cultural, and psychological factors. These associations are open-ended and not content-like: they vary across speakers, time and dialects, and for a given speaker across different times and circumstances. However, while variable, they share a common thread, insofar as they are rooted in the same set of socio-historical and cultural facts. Finally, and importantly, these associations do not attach to words, but rather to their articulations: (certain of the) slurs’ orthographic and phonological forms.

Our account—the Articulation Account—is uniquely well positioned to capture the full range of data concerning the slurs’ pejorative sting. First, it captures that a denial of any specific pejorative content alongside the predication of a slur doesn’t lead to *linguistic* incoher-

⁸ See, for instance, Heller 2000.

⁹ <https://www.cbc.ca/news/canada/toy-pandas-bearing-swastikas-a-cultural-mix-up-1.343550>; <https://www.theglobeandmail.com/news/national/swastikas-still-hurt-toy-error-shows/article4143067/>.

ence: slurs do not semantically encode, or pragmatically convey, any such content, and so, there's no inconsistency or incoherence in predicating a slur, while denying such content. The account can also explain why there can be variation in the presence and severity of the offensive sting across time and dialects, as well as across different speakers, and even for a single speaker at different times, without positing a change in meaning, or linguistic confusion. The associations that tie to articulations change dynamically over time, and across communities, as a result of changes in various socio-historical and cultural factors; and they are open-ended, in that no particular bit of content or affect is semantically encoded, nor it must be evoked by any given agent in order for them to be linguistically competent.

The Articulation Account also explains why the pejorative sting is hyperprojective, indeed, why it persists even in purely quotative environments and mere displays of slurs: for it is the articulation of the slur, which is present in such environments, that carries the negative associations that constitute the pejorative sting.

Similarly, we can explain why the tokening of a slur is neither necessary nor sufficient for triggering the pejorative sting. On the one hand, the articulation itself can trigger the pejorative sting without the slur being tokened. This is already reflected in the phenomenon of inheritance: in tokening the Mandarin demonstrative, for instance, the speaker didn't thereby token a slur. It is the articulation of the slur that's tokened, due to the orthographic—or rather, phonological—accident, which in turn carries the pejorative sting. That no slur is tokened in this case should be evident as soon as one reflects on the fact that the written counterpart, '那个', triggers no pejorative sting whatsoever. To insist otherwise would be to confuse words and their articulations.¹⁰

On the other hand, absent a slurring articulation a slur loses its pejorative sting: if one attempts to token a slur, but misarticulates it so badly that it's difficult to even recognize, the tokening of the slur loses its pejorative sting. If one knows the speaker misarticulated a slur, one might then wonder whether the speaker had malicious intentions or not, whether they are blameworthy or not, and so forth, but one won't thereby experience the same viscerally palpable effect as if one had been confronted by the articulation of the slur directly. This is just as if someone tries to punch you, but misses, or barely grazes you. If you know what they tried to do, you might wonder whether they had any malicious intentions, whether they are blameworthy and so forth; however, you won't thereby experience the pain of the would-be punch.

One might be tempted by the thought that we've conflated metaphysics and epistemology: isn't it simply the case that absent (standard) slurring articulations the pejorative sting is absent because we

¹⁰ For further discussion of the important distinction between words and articulations, see, e.g., Stojnić 2001; Hawthorne & Lepore 2011; Kaplan 1990; 2011.

fail to recognize that the word tokened is a slur? And conversely, isn't it simply because we confuse standard articulations of slurs for the tokenings of slurs, even when no slur is tokened, that such articulations retain the pejorative sting? If so, this is consistent with the idea that it's the slurs that carry the pejorative sting, while the articulations merely serve as evidence as to whether a slur was tokened or not.

This reaction should be resisted. Note that in none of the inheritance examples above there's any confusion about which word was tokened. For instance, the speaker in the Mandarin demonstrative incident announces ahead of time which word of which language will be tokened. So, these examples cannot be chalked up to mere confusion. Nor should we think that the absence of the offensive sting in the absence of (certain of) the standard articulations of slurs is merely a byproduct of epistemic confusion, a failure to recognize that a slur has been tokened.

To see this, it is useful to reflect on another important phenomenon that has been largely neglected in the literature—the variance in the presence and severity of the offensive effect across different articulations of one and the same slur term. A vivid example of this is the phenomenon of graphic pejoratives in logographic languages, like Mandarin. Graphic pejoratives allow that only certain written, but not spoken, articulations of a particular term can carry the offensive sting. For example, historically, many exonyms were considered as pejorative only when written with one of several possible, phonologically indistinguishable phono-semantic compounds (characters). Matisoff reports that names of people deemed lesser would be deemed offensive when they featured the 'beast radical' (豸) (e.g., '獠,' i.e., 'Yao'), a rendering typically avoided in formal correspondences (Matisoff, 1986: 6). Thus, we have a case where only one of several possible standard articulations of a single term, a particular spelling, carries the pejorative sting. If the sting were a matter of a word or its meaning, this would be baffling, for a word retains its identity, and meaning, regardless of the choice of whatever acceptable spelling. Nor can it be that articulations of slurs trigger offense merely derivatively, because we recognize them as articulations of slurs, and that when slurs are tokened absent slurring articulations, the sting is absent due to a failure to recognize that a slur is tokened. With graphic pejoratives, standard articulations of the same word vary in whether they trigger the offensive sting. Yet, the word is what it is, and means what it means, regardless of how you spell or pronounce it.

The phenomenon of variance in the offensive sting across different articulations of one and the same term is likewise not isolated, nor exclusive to logographic languages. In *Inflammatory Language* we discuss a wide range of such cases, including the use of minced oaths, variations in spelling or pronunciation designed to blunt the offensive sting, as well as similar uses of the censoring asterisk, as in:

8.
 - a. Hermione is a mudbl**d.
 - b. The word ‘mudbl**d’ is censored in this sentence.

Indeed, variation in the associations attached to different articulations of a single term can explain the effects of variation in articulation that sometimes accompanies reclamation—the reclaimed in-group uses of a slur as a term of pride, a form of solidarity, camaraderie, activism, or simply a neutral moniker of self-designation. Consider, for instance, the in-group uses of the N-word in its non-rhotic pronunciation (and the corresponding spelling variation). While out-group tokenings of both non-rhotic and rhotic articulations are unequivocally offensive, there is a debate both over the availability for in-group reclamation, and comparative offensiveness, of these different articulations of the slur. Whatever one makes of such debates, they wouldn’t even *prima facie* make sense if it were the term itself that already triggered the pejorative sting. On the Articulation Account, by contrast, the phenomenon is entirely unremarkable: different articulations of one and the same term can vary in associations and thus in the presence and severity of the offensive sting.

In sum, we’ve argued that the pejorative effect of slurs is constituted by a cluster of associations that tie to (certain of) the *articulations* of slurs, due to a variety of socio-historical, cultural, psychological, and other factors. While it might be *prima facie* unpalatable to accept that the pejorative sting of slurs is not a matter of meaning, or even language, we’ve argued that there’s overwhelming evidence for this view. The Articulation Account is uniquely well-positioned to explain the full range of puzzling data concerning the behavior of slurs and their pejorative sting.

We emphasize that while associations that constitute the pejorative sting can be weaponized to denigrate, harm or offend, they can also be triggered accidentally or negligently, or exploited for a variety of other effects: rhetorical, pedagogical, artistic and so on. But such tokenings do remain inherently risky. Whether any particular tokening is apt and necessary, or problematic, flippant or gratuitous, whether it constitutes a transgression, let alone one with special moral timbre, and whether any offense it causes is warranted, these are all questions parasitic on the fact that the tokenings of slurring articulations trigger the pejorative sting. But ultimately, we argue, the sting is not a matter of semantics or pragmatics, or even language: it is not slurring *words*, but *articulations* that carry the offensive sting.

References

- Anderson, L., and Lepore, E. 2013. “Slurring Words.” *Noûs* 47 (1): 25–48.
- Brandom, R. 2002. *Articulating Reasons: An Introduction to Inferentialism*. Cambridge: Harvard University Press.

- Camp, E. 2013. "Slurring Perspectives." *Analytic Philosophy* 54 (3): 330–349.
- Carnaghi, A., and Mass, A. 2008. "Derogatory language in intergroup context: Are 'gay' and 'fag' synonymous?" In Y. Kashima, K. Fiedler, & P. Freytag (ed.). *Language-based approaches to the formation, maintenance, and transformation of stereotypes*, 117–134.
- Cepollaro, B. 2015. "In Defense of a Presuppositional Account of Slurs." *Language Sciences* 52: 36–45.
- Cepollaro, B., and Stojanovic, I. 2016. "Hybrid evaluatives." *Grazer Philosophische Studien* 93 (3): 458–488.
- Cepollaro, B., Sulpizio, S., and Bianchi, C. 2019. "How bad is it to report a slur? An empirical investigation." *Journal of Pragmatics* 146: 32–42.
- Dummett, M. 1981. *Frege: Philosophy of Language*. Cambridge: Harvard University Press.
- Fasoli, F., Paladino, M. P., Carnaghi, A., Jetten, J., Brock, B., and Bain, P. G. 2015. "Not 'just words': Exposure to homophobic epithets leads to dehumanizing and physical distancing from gay men." *European Journal of Social Psychology* 46 (2): 237–248.
- Frege, G. 1897/1979. "Logic." In H. Hermes, F. Kambartel, & K. F. (ed.). *Posthumous Writings: Gottlob Frege*. Blackwell, 126–151.
- Hawthorne, J., and Lepore, E. 2011. "On Words." *Journal of Philosophy* 108 (9): 447–485.
- Heller, S. 2000. *The Swastika: Symbol Beyond Redemption?*. New York: Allworth Press.
- Hom, C. 2008. "The semantics of racial epithets." *Journal of Philosophy* 105 (8): 416–440.
- Hom, C., and May, R. 2013. "Moral and semantic innocence." *Analytic Philosophy* 54 (3): 293–313.
- Jeshion, R. 2013. "Expressivism and the offensiveness of slurs." *Philosophical Perspectives* 27 (1): 231–259.
- Jorgensen Bolinger, R. 2017. "Pragmatics of Slurs." *Noûs* 51 (3): 439–462.
- Kaplan, D. 1990. "Words." *Aristotelian Society Supplementary Volume* 64 (1): 93–119.
- Kaplan, D. 2011. "Words on Words." *Journal of Philosophy* 108 (9): 504–529.
- Kennedy, R. 2002. *Nigger: The Strange Career of a Troublesome Word*. New York: Pantheon Books.
- Lepore, E., and Stone, M. 2018. "The Pejorative Tone." In D. Sosa (ed.). *Bad Words: Philosophical Perspectives on Slurs*. Oxford: Oxford University Press, 132–154.
- Matisoff, J. 1986. "The Languages and Dialects of Tibeto-Burman: An Alphabetic/Genetic Listing, with Some Prefatory Remarks on Ethnonymic and Glossonymic Complications." In J. McCoy, Light, & T (ed.). *Contribution to Sino-Tibetan Studies*. Brill, 3–57.
- McCready, E. 2013. "Varieties of conventional implicature." *Semantics and Pragmatics* 3: 1–5.
- Neufeld, E. 2019. "An essentialist theory of the meaning of slurs." *Philosophers' Imprint* 19: 1–29.

- Nunberg, G. 2018. "The Social Life of Slurs." In N. W. Acts, Godal, Daniel; Harris, Daniel; Moss, Matt (ed.). *New Work on Speech Acts*. Oxford: Oxford, 237–293.
- O'Hehir, A. 2020. "So much for youth apathy: Student radicalism escapes the '60s at last." Retrieved from *Salon*: https://www.salon.com/2015/11/17/so_much_for_youth_apathy_student_radicalism_escapes_the_60s_at_last/
- Potts, C. 2007. "The expressive dimension." *Theoretical Linguistics* 33 (2): 165–198.
- Quine, W. 1953. *From a Logical Point of View*. Cambridge: Harvard University Press.
- Schlenker, P. 2007. "Expressive presuppositions." *Theoretical Linguistics* 33 (2): 237–245.
- Stojnić, U. 2022. "Just Words: Intentions, Tolerance, and Lexical Selection." *Philosophy and Phenomenological Research* 105 (1): 3–17.
- Stojnić, U., & Lepore, E. 2025. *Inflammatory Language: Its Linguistics and Philosophy*. Oxford: Oxford University Press.
- Williamson, T. 2009. "Reference, inference, and the semantics of pejoratives." In L. P. Almog Joseph (ed.). *The Philosophy of David Kaplan*. Oxford: Oxford University, 137–158.