



Doxastic and Epistemic Sources of Offense for Slurring Terms

David Miguel Gray^{1,2}

Received: 7 August 2024 / Accepted: 30 August 2024
© The Author(s) 2024

Abstract

Existing analyses of slurs emphasize how linguistic mechanisms make slurs derogatory. I will argue that, in addition to linguistic mechanisms, there are overlooked doxastic and epistemic features of standard slurring utterances that can be sources of offense. Additionally, I argue that the doxastic feature that distinguishes slurring utterances from other negatively valenced utterances is fundamental to understanding slurring terms. Clinical Trial Registration: Not applicable.

Keywords Slurs · Racial epithets · Doxastic · Epistemic · Offense · Pluralism · Derogatory

1 Introduction

Slurring utterances are reasonably taken to be offensive. Philosophers and linguists have spent considerable time giving arguments as to what particular feature or mechanism of a slurring utterance acts as the source of offense. While virtually all accounts of slurs have claimed that the mechanism is linguistic,¹ I will argue

¹ Posited linguistic mechanisms include phonetic features (Mandelbaum et al., 2024), truth-conditional content (Bach, 2018; Hom, 2008; Neufeld, 2019; Sennet & Copp, 2015), the expressive attitudinal component (Boisvert, 2008; Croom, 2011; ‘Expressives and beyond’, 2013; Jeshion, 2013; McCready, 2010; Potts, 2007; Richard, 2008; Saka, 2007), conventional implicature (Lycan, 2015; McCready, 2010; Whiting, 2013; Williamson, 2009), presuppositional content (Cepollaro, 2015; Cepollaro & Stojanovic, 2016; Predelli, 2010; Schlenker, 2007), inferential role (Dummett, 1973; Tirrell, 1999), meta-content (Blakemore, 2015), perspective signaling (Camp, 2013, 2018), conversational implicature (Bolinger, 2017; Nunberg, 2018), and the sociolinguistic mechanism of prohibition/taboo (Anderson & Lepore, 2013).

✉ David Miguel Gray
dmgray2@memphis.edu; davidmiguelgray@gmail.com

¹ Philosophy Department, The University of Memphis, Memphis, USA

² Institute for Intelligent Systems, Fed Ex Institute of Technology, The University of Memphis, Memphis, USA

that—in addition to whatever linguistic mechanisms derogate²—there are doxastic and epistemic sources of offense. I will arrive at this point by showing that slurring utterances involve a kind of doxastic error and that this error is fundamental to understanding the category of slurs. This error also explains at least one of the sources of offense common to slurring utterances.³ I take my arguments for a doxastic and epistemic source of offense, as well as the method I use to say what a slur is, to be original contributions to the literature on slurs. Additionally, in presenting this argument and allowing for multiple sources of offense, I am in opposition to ‘monistic’ accounts of slurs that claim that there is only a single source of offense. Most linguistic accounts of offense fall into this category. Jeshion (2013), Popa-Wyatt and Wyatt (2018), and Mandelbaum et al. (2024) are exceptions in acknowledging multiple sources of offense (but silent regarding non-linguistic sources).

I will define some terminology to clarify this discussion. Given the variety of negative practices utterances can serve, I will use the phrase ‘negatively valenced’ to talk about these utterances regardless of whether they use slurring terms. I use ‘negatively valenced’ broadly to indicate phrases or terms used to say something bad about the person or group being spoken about—and what is said could be either appropriate or inappropriate. Utterances that belittle, blame, deride, derogate, rebuke, and reprimand are negatively valenced. A ‘negatively valenced term’ is a term whose normal usage makes the utterance negatively valenced. For instance, sincerely calling someone a ‘thief,’ a ‘plunderer,’ or a ‘usurper’ are negatively valenced terms. Slurs and some insults are also negatively valenced terms. The ‘target’ of a slur may include groups or members of a group based on race, gender, nationality, migratory status, and mental or physical ability.⁴ A ‘standard use’ or ‘targeted use’ of a slur involves an utterance of a slur by a competent language user aiming to pick out the group or its members commonly associated with the slur. Standard use cases are normally associated with negative intent, such as the intent to insult.⁵ I do not have an exhaustive list

² Since extant accounts of slurs claim that the source(s) of offense is linguistic, this source of offense is often referred to as the ‘derogatory force’ of the slurring utterance. ‘Derogatory force’ is a term of art for the derisive powers of utterances that stem from a locutionary (semantic) act or an illocutionary (pragmatic) act. Utterances aside from those using slurs can derogate, and in some cases, offense would be warranted (e.g., calling a murderer a ‘murderer’). Since normal usage of slurs do warrant offense, I will use the phrase ‘sources of offense’ to cover all speaker-side offense-causing properties of negatively-valenced utterances, regardless of the nature of the mechanism—linguistic, doxastic, or otherwise. These speaker-side mechanisms can be contrasted with the hearer-side phenomenon of actually taking offense by virtue of the speaker’s utterance. This phenomenon is often explained as a perlocutionary act on the part of the speaker, e.g., Austin (1962, 101), and a perlocutionary effect on the part of the audience. This kind of offense is partially dependent on the hearer’s attitudes and thus will not be a central concern here.

³ Additionally, I take it that doxastic and epistemic features being sources of offense does not settle any issues regarding which linguistic mechanism(s) are sources of offense.

⁴ For example, see (Anderson & Lepore, 2013; Cepollaro & Stojanovic, 2016; DiFranco, 2017; Gray & Lennertz, 2020; Hom, 2008; Jeshion, 2013; Neufeld, 2019; Nunberg, 2018; Spotorno & Bianchi, 2015).

⁵ Bach and Harnish (1979: 17) note that the characterization of perlocutionary acts given by Austin (which I characterized in fn. 2) is not very useful as any utterance can have a variety of unintended effects on those they are directed at. While Austin leaves it open whether the speech act needs to have the intent to produce the desired effect in the audience, Bach and Harnish usefully suggest limiting the idea of perlocutionary acts just to those that have the intent to produce a particular effect on the audience. Following this suggestion, we have an alternative characterization of a targeted or standard use of a slur as one that intends to offend.

of non-standard use cases, but some examples are utterances of the slur when it is embedded in quotation (pure or direct quotation), many cases of indirect quotation, possibly corrective and pedagogical uses (Hornsby, 2001: 128–29; Hom, 2008: 429), and some instances of negation.⁶ Finally, I will follow Bollinger (2017) in using the phrase ‘warranted offense’ to refer to an offense that one would be justified in taking, regardless of whether one is actually offended. This type of offense does not need to have the affective or physiological component usually associated with actual offense. I take warranted offense to require a reason to be offended. Being targeted by a slur or hearing someone target others with a slur warrants offense. When an offense is warranted, it then makes sense to talk about the source of the offense. With that terminology in place, I can now state my thesis more precisely. I will argue that targeted uses of slurring terms have doxastic and epistemic sources of warranted offense. Additionally, the nature of this doxastic offense will help distinguish slurs from other negatively valenced terms.

In §2, I will start with an analysis of negatively valenced utterances and provide a brief taxonomy that allows us to distinguish slurring utterances from other negatively valenced utterances. Some negatively valenced utterances do not use negatively valenced terms (§2.1), and some negatively valenced utterances use negatively valenced terms that are not slurring terms (§2.2.1, §2.2.2). Explaining what are and are not slurring utterances and terms is crucial to an understanding of slurs (§2.2.3). I will argue that standard slur utterances fundamentally involve an evaluative error that distinguishes slurs from other negatively valenced terms. I will argue in §3 that in many targeted uses of slurring terms, there is a doxastic source of offense (§3.1) and an epistemic source of offense (§3.2). Both sources of offense can cause warranted offense. I will close in §4 by responding to some potential objections in addition to motivating pluralism with respect to sources of offense for slurring terms.

2 Locating Slurs

2.1 Negatively Valenced Utterances without Negatively Valenced Terms

I will be concerned with negatively valenced utterances that target groups and individuals by virtue of being members of groups. An utterance can be negatively valenced without using a group term that is negatively valenced (e.g., ‘damn lawyers’ or an exasperated utterance of ‘teenagers!’). I only mention negatively valenced utterances without negatively valenced terms to set them aside. My focus will be on negatively valenced utterances that use negatively valenced terms.

⁶ This characterization of standard usage will help to demarcate cases of slur utterances that involve doxastic errors, to be discussed in Section 3.1

2.2 Negatively Valenced Utterances with Negatively Valenced Terms

Negatively valenced terms include ‘thief,’ ‘killer,’ ‘asshole,’ ‘Karen,’ ‘spic,’ and ‘honkey.’ We can divide these terms into three subgroups. These three subgroups are distinguishable by how they answer the following question. When can these negatively valenced terms be applied to a group or its members in virtue of a group property that is ‘morally relevant’? By ‘morally relevant,’ I mean a property that warrants a positive or negative evaluation.

2.2.1 Negatively Valenced Terms that can Pick out Morally Relevant Properties

This subgroup contains terms such as ‘thief’ or ‘killer.’ These terms can derogate. The derogation can be warranted in cases where someone steals or kills without good reason. In such cases, to call someone a ‘thief’ or ‘killer’ is to derogate for morally appropriate reasons. In these cases, offense on the part of the targets is not warranted. As such, these negatively valenced terms, when used appropriately, are not sources of warranted offense.

While stealing and killing can be morally relevant properties and are normally seen as wrong, stealing and killing can also be seen as not worthy of negative evaluation in certain cases. Similarly, the terms ‘thief’ and ‘killer’ are negatively valenced because it is, *prima facie*, bad to be to be a thief or a killer. Nevertheless, the negative valence of these terms is ‘cancellable.’⁷ This is the case with ‘thief’ or ‘killer’:

- (1) Theodore is a thief, but he only stole food for his family in a time of great poverty.
- (2) Kim is a killer, but that is, unfortunately, the expectation of combat soldiers in just wars.

I am claiming that, even when correctly applied to people who have stolen or killed, the negatively valenced terms ‘thief’ and ‘killer’ need not pick out morally relevant properties. However, they often do.

How these negatively valenced terms can warrant offense is worth considering for the upcoming discussion on slurs. One such case would be where the terms ‘thief’ and ‘killer’ are applied in error. Suppose Theodore and Kim did not steal or kill, respectively. In that case, targeting them with the terms ‘thief’ and ‘killer’ is morally inappropriate: they have not performed the required actions to make evaluating them with these terms morally apt. In such a case, Theodore and Kim are warranted in taking offense. To not put too fine a point on it—as an analysis of these terms is not my main objective—we can say what warrants the offense has more to do with the mistaken evaluation on the part of the speaker.

⁷ I am not taking a stand here on whether the negativity of these utterances results from conversational implicature or some other linguistic mechanism. While the reliability of the cancellability test for conversational implicature has come under fire, it is important to remember that for Grice, cancellability is only a necessary condition for conversational implicature (Grice, 1989: 44).

2.2.2 Negatively Valenced Terms that do Pick out Morally Relevant Properties

This second subgroup includes insults like ‘asshole’ but also terms like ‘Karen.’ I acknowledge that there can be playful uses of ‘asshole’ or even ‘Karen.’ However, I take these to be akin to appropriated uses of slurs in that they are polysemous but not negatively valenced. (e.g., see Gray & Lennertz, 2020). Here, I limit myself to the sincere use of the negatively valenced terms. ‘Asshole’ targets people whose behavior we disapprove of. For example, a speaker who holds an evaluative belief—let us say, the belief that Ashley is an asshole—is justified in that belief if she also believes that Ashley has properties that are morally relevant to being an asshole. If I have a good reason for believing that Ashley mistreats her friends for laughs, this would justify my evaluative belief that Ashley is an asshole.

Unlike the previous category, which also targets morally relevant properties, these negatively valenced terms are not cancellable. While someone could say, ‘Yes, Ashley is an asshole, but you have to be when all of your competitors are even bigger assholes.’ this does not cancel the term’s negative valence—it just contrasts bad behavior with even worse behavior.

My focus is on negatively valenced terms that can target individuals—as members of groups—and groups. We can talk of the group ‘assholes,’ but it is admittedly forced. However, other terms in this category include terms like ‘Karen.’ (This term is commonly confused with a slur, but I will argue it is fundamentally different.) Karens have been called ‘the policewomen of all human behavior’ (Miller, 2019). In more detail, the term ‘Karen’ picks out white women who have a lack of awareness of social hierarchies, are demanding, have a strong sense of entitlement, and regularly act on this entitlement (Lewis, 2020). The occurrence of such behavior has been noticed with such regularity in recent times that individuals exhibiting these properties are discussed as a group. This is undoubtedly due to the prevalence of problematic racial and customer-service interactions popularised in the media. It is the morally relevant property of the group members (viz., entitled and demanding white women who attempt to exploit a type of racial privilege) that leads to their derogation.

What makes someone a member of the group Karens or assholes is a morally relevant property. Additionally, in this subgroup, the morally relevant property is the basis for group membership. Whereas ‘thief’ and even ‘killer’ can be defined using value-neutral language, ‘Karen’ and ‘asshole’ are not—at least not without warranting a critique of euphemism.⁸ As such, the negative valence of terms in this group is not cancellable.

⁸ Contrast the more neutral definitions for ‘thief’ as ‘One who takes portable property from another without the knowledge or consent of the latter, converting it to his own use’ with the less neutral ‘asshole,’ ‘A stupid, irritating, or contemptible person; a person who behaves despicably’ and ‘Karen,’ ‘An obnoxious, angry, entitled, and often racist middle-aged white woman who uses her privilege to get her way or police other people’s behaviors’ (Thief, n. 2022; Asshole, n. 2022; Karen, n. 2020).

2.2.3 Negatively Valenced Terms that never Pick out Morally Relevant Properties

Returning to our question of whether negatively valenced terms can be applied to a group or its members by virtue of a group property that is morally relevant, our third subgroup will answer in the negative. I take slurs to occupy this category.

Targeted slur utterances negatively evaluate groups or members of a group *by virtue of* being members of that group. If this kind of evaluation were evaluatively *apt*, then the properties that serve as the basis for group membership would have to be morally relevant. (By ‘evaluatively apt,’ I mean the group or its members would have to have a morally relevant group property suitable for positive or negative evaluation.) However, none of the group properties are morally relevant.⁹ Therefore, targeting groups with slurring utterances is never apt, and the belief that the targets of slurring terms are evaluatively apt is categorically false.

This feature marks a difference between slurs and the negatively valenced terms captured by the other two categories. We can explain this difference in two ways. On the one hand, terms like ‘thief,’ ‘killer,’ ‘asshole,’ and ‘Karen’ can be evaluatively apt when the relevant group property is present. This cannot be the case for slurs since there is no relevant group property that can be subject to moral evaluation. On the other hand, negatively valenced terms in all three categories can be misapplied. However, the nature of the error is different. In the case of evaluating someone with the terms ‘thief’ or ‘killer,’ I might mistakenly think that someone stole or killed when they did not, or I might not realize that there were extenuating circumstances that made the acts permissible. In the case of evaluating someone with the terms ‘asshole’ or ‘Karen,’ I might misinterpret a behavior by not recognizing some important contextual features (say someone was acting like an asshole because they were practicing a part for a play). However, in the case of evaluating someone with a slur, the error does not lie in thinking someone has negative morally relevant properties that someone *could* have but does not. Rather, the error is assigning negative morally relevant group properties that no one or group could have. Targeting with a slur amounts to a kind of moral category mistake.¹⁰ In this respect, targeted uses of slurring terms fundamentally involve a doxastic error.

⁹ I should say sometimes words that start as slurs can develop different uses, making them polysemous with the slur. In Chris Rock’s famous 1996 comedy special *Bring the Pain*, he distinguishes between Blacks and N*gg*s in such a way that the extension of the latter group narrows to Blacks that exhibit a negative stereotypical set of morally relevant properties (Truesdell, 1996). One interpretation of how this occurs is that the stereotypical properties implied by its use as a slur become part of its lexical meaning (Nunberg, 2018).

¹⁰ One difference between a speaker incorrectly calling me a ‘thief’ and targeting me with a slur is that—when called a thief—I am more likely to correct the speaker. However, shouldn’t I be equally inclined to correct the speaker using a slur, given that I claim both involve an error? On the one hand, correction of slurring uses is not uncommon in philosophy (e.g., see (Hornsby, 2001: 128–29; Hom, 2008: 429) or outside of philosophy (e.g., see Fletcher 1997; Ebony, 1982, 122). On the other hand, I think correction might more readily come to mind for terms like ‘thief’ as it can both correctly or erroneously be applied. Additionally, given the vitriol that often goes with a slurring utterance, concerns such as safety may override the target’s wanting to correct the speaker.

In this section, I have developed a taxonomy for individuating slurs that is largely non-linguistic. Rather, by using the ideas of doxastic error, morally relevant properties, and evaluative inaptness, my taxonomy makes moral and epistemic considerations central to individuating slurs. One advantage of this is that the features that allow us to categorize slurs as distinct from other negatively valenced terms also reveal what makes standard uses of slurs wrong.

One could argue that there is no need for this alternative taxonomy. This is because a way of distinguishing slurs from other negatively valenced terms already exists. Many claim that what makes slurs distinctive is that they have a neutral (non-evaluative) counterpart that is descriptively detachable (Hornsby, 2001: 129; Schlenker, 2007: 238; Whiting, 2013: 364). For instance, a slur like ‘honkey’ has the neutral counterpart ‘white.’¹¹ However, terms like ‘asshole’ or ‘Karen’ do not have neutral counterparts.

One response we can give is that given by Lauren Ashwell (2016). Ashwell convincingly argues that some negatively valenced terms, such as ‘slut’, are slurs and do not have a neutral counterpart. For instance, “a woman who has too many sexual partners” is still evaluative since what counts as ‘too many’—or even ‘many’—is still evaluative (Ashwell, 2016: 234–35). Similarly, Jeshion (2017: 135) takes slurs like ‘Hillbilly,’ ‘Redneck,’ and ‘White Trash’ to not have neutral correlates. It is worth noting that Ashwell and Jeshion are making different points. Jeshion recognizes, as a matter of the factual record, that there are no nonpejorative corollary *terms* for some slurs. Ashwell argues that, in some cases where there is no nonpejorative corollary term, there is also no adequate corollary *description*. The reason is that some descriptions, like that from Ashwell quoted above, cannot remove the negative valence of the slur while keeping intact a description that would still be co-extensive with the intended referent of the slur. Unlike a taxonomy based on the idea of descriptively detachable neutral counterparts, my view requires neither a neutral corollary term nor a description for individuating slurs—just the evaluation of non-evaluatively apt group properties. As such, my taxonomy avoids both of these issues.

3 Doxastic and Epistemic Sources of Offense

3.1 A Doxastic Source of Offense

The belief that slurs are evaluatively apt could be explicitly held by the speaker, but it need not be.¹² To utter a slur, in standard-use cases, requires the tacit belief that

¹¹ This is also referred to as the ‘nonpejorative correlate’ of a slur (Hom, 2008: 417).

¹² A major motivation for providing an account of tacit beliefs is that they are used to explain action in cases where action requires a particular belief, and requiring that that belief always be explicit is too demanding or unrealistic. I am following that practice here. However, there could be many cases where the content of the belief I am claiming is a hallmark feature of standard-use cases of slurs is explicitly held. If we take it, as I do, that a tacit belief that p and an explicit belief that p for a given subject are mutually exclusive, my claim should be understood disjunctively as requiring either an explicit or tacit belief that using a slur to target a group or member of a group is evaluatively apt.

the target is evaluatively apt by virtue of evaluable group properties.¹³ (There are non-standard use cases that don't involve There are, of course, many other beliefs that such a speaker might also hold, explicitly or tacitly, that could also be offensive (e.g., that minorities are not deserving of protection from hate speech, that they are thieving, etc.). What makes the doxastic error I have proposed central to standard-use cases of slurring utterances is that it captures what the speaker is committed to in standard-use cases. That is, a slurring utterance negatively evaluates its target. Additionally, if we take slurs to be inappropriate, unjust, or wrong, the evaluation is unwarranted. The categorically false nature of the doxastic error I have suggested both demarcates the category of slurring terms and explains one reason that they are offensive.

It could be suggested that what I am claiming is a doxastic source of offense is actually a linguistic feature of slurring terms. For instance, one might think that slurring utterances have the presuppositional content that the target of the slur is apt for evaluation by virtue of the member's (or groups') group properties. This claim could be true, as far as it goes. However, there are still problems with this suggestion. Extant accounts that claim presuppositional content is a source of offense make very different claims than mine. For example, Schlenker takes it that slurs are expressives that have a presupposition which is, in part, "attitudinal (it predicates something of the mental state of the agent of that context)" (2007: 237). For example, the source of offense for the slur 'honkey' in a context *c* in a world *w* would be the agent of *c* believing "in the world of *c* that white people are despicable" (Schlenker, 2007: 238). Predelli (2010) and Cepollaro (2015) use a similar approach, which can be captured by the simpler formulation: 'the 'honkey'-users despise whites for being whites.' Both of these accounts give a central, albeit different, role for the source of offense: the presuppositional attitude of despising. In contrast, my claim about the speaker's tacit belief is not offensive by virtue of an attitudinal feature.

What I am claiming is that it is the falsity of the tacit, or explicit, belief that an individual or group is morally apt for a negative evaluation—by virtue of a non-evaluatively apt group property—is a doxastic source of offense. I take this source of offense to be present in the first two subcategories of negatively valenced terms in cases where the use of negatively valenced terms mischaracterizes their targets. Additionally, as I discussed regarding the misapplication of the term 'thief,' the offense to be taken was not a feature of the term 'thief' as the warranted offense only occurs when the term is misapplied.¹⁴ This suggests that in the case of slurs, even if there are linguistic mechanisms that are sources of offense, this is no reason to hold that there cannot be non-linguistic doxastic and epistemic sources of offense.

¹³ In characterizing the speaker's false belief as tacit, I have in mind something similar to Crimmins (1992: 248–249). I would say that a speaker *S* who utters a slur tacitly believes that *p* (i.e., groups, or their members, are evaluatively apt in virtue of group properties) just in case *S*'s cognitive dispositions are relevantly as if *S* has an explicit belief in *p*.

¹⁴ Additionally, in the case of slurs, the fact that the error is categorically false may add to their insidious nature.

3.2 An Epistemic Source of Offense

The idea of mischaracterization and negative evaluation leads us to an epistemic source of offense (albeit less perfectly, as even tacit false beliefs can be justified). For instance, if Theodore is called a thief incorrectly, the speaker has acted as if Theodore has stolen without good reason. In addition to the offense warranted by the mischaracterization, how the speaker came to such an evaluation of Theodore could also offend. If it was the case that the speaker has jumped to a conclusion—that her evaluation is supported by suspect hearsay or other questionable evidence—it displays a lack of epistemic care. Additionally, given that the doxastic error is negatively evaluative, the degree of epistemic scrutiny ought to be higher than for beliefs that lead to non-evaluative claims. This is because (a) one is not obligated to make such an evaluation, even if it is permissible, and (b) a misevaluation wrongs and can harm a person. Epistemic due diligence is required in cases of negative evaluations: one must exercise care in collecting, weighing, and vetting evidence to avoid harming others.

Returning to the case of targeted uses of slurring terms, we can similarly question the justificatory status of beliefs that support the negative evaluation. One can imagine a case in which a person is reasonably justified in their evaluative belief. If one was embedded in an insulated community that espoused a racist ideology—claiming that races had innate morally reprehensible properties—and had no evidence to the contrary, the doxastic error would be justified. In this case, there is still a doxastic source of offense, but the justificatory fault is not that of the speaker. Offense can still be warranted in this case, but the epistemic source will not lie with the speaker but with the community that acts as a testimonial source for the speaker.

The case just imagined is a rarity. In most cases where slurs are uttered, the shoddy or absent effort of subjecting potential evaluations to epistemic scrutiny can itself be a source of offense.

4 Further Considerations

In contrast to accounts of slurs that claim the sources of offense are linguistic (broadly construed to include pragmatic and sociolinguistic mechanisms), I have argued that there are also important doxastic and epistemic sources of offense. Additionally, the doxastic feature of slurring terms I have proposed both provides a method for demarcating what slurs are and suggests one of the things that makes slurring wrong.

My view is pluralistic. I take there to be epistemic, doxastic, and linguistic sources of offense. Slur utterances, like much of speech, involve speaker intentions, word choice, and a variety of semantic and pragmatic mechanisms. I take it that several or all of these mechanisms could produce offense.

I should note that doxastic and epistemic sources of offense may not accompany non-standard slur utterances, such as direct and indirect quotations, as well as educational contexts where one does not hold that false offensive belief, explicitly or tacitly. This could be seen as making my project less interesting, as there are monistic theories that claim they can explain all the relevant phenomena. There are three reasons to think this is not the case.

First, that a monistic theory can explain all the features of slurs is currently a promissory note. This should be no surprise, as there is little consensus in many areas of philosophy. Second, a monistic theory of slurs is preferable only under the assumption that there can be only one mechanism, which explains why slur utterances warrant offense. Two potential errors are being made here. While I have suggested a way to think of slurs as a singular category—and that this way of thinking about them explains a source of offense—this does not mean that there cannot be other sources of offense. Perhaps uttering slurs violates a sociolinguistic norm of polite conversation, and that offends. Perhaps slurring utterances come with a negative attitude that can offend. Perhaps the at-issue content of a slur is also offensive. These possibilities lead to the second potential error. In non-slurring cases, it seems we can think of a variety of linguistic mechanisms that could cause offense. What one literally says or what one implies by what is said could both be sources of offense in the right circumstances. The defender of a monistic account for slurs not only owes a story of how one source of offense can explain all the phenomena we associate with a slur's offensiveness but also why that mechanism blocks other mechanisms from being offensive just in the case of slur utterances. Without answers to these questions, we cannot assume a monistic approach to slurs is better than a pluralist one.

Another reason to consider pluralism is that some features of slurs might be better explained on a pluralist account. For instance, several philosophers take it that the same slur, during the same period, can vary in how offensive it is depending on how it is used (see Bolinger, 2017; Cepollaro et al., 2019; Nunberg, 2018; Popa-Wyatt & Wyatt, 2018). For instance, we might think that standard uses of slurs warrant greater offense than cases of indirect quotation. If this were the case, then any theory of slurs would owe an account of this variation.¹⁵ Empirical evidence has been provided for the claim that a slur's offensiveness varies in different use cases (Cepollaro et al., 2019). In their study, 111 subjects were presented with direct and indirect sentences that predicated slurs to individuals. The direct sentences had the form "Y: X is a P," and the indirect sentences had the form "Z: Y said that X is a P." X, Y, and Z were replaced with proper names (roughly half male and half female names) and the predicates were slurring terms (2019, pp. 37–38). Their study showed that slurs were rated significantly more offensive in the direct speech condition and less offensive in the indirect speech condition.¹⁶ The authors claim that certain monistic positions will have trouble explaining this data, such as Prohibitionism (e.g., see Anderson & Lepore, 2013)—which claims that all non-appropriated uses of slurs are offensive in virtue of violating sociolinguistic prohibitions on their utterance.¹⁷

¹⁵ Interestingly, very few theorists of slurs seem concerned with this phenomenon. Instead, there is a much greater focus on how different slurs can vary in how offensive they are in the same time period (e.g., Anderson & Lepore, 2013; Bolinger, 2017; Camp, 2013; Croom, 2011; Čupković, 2015; Hom, 2008, 2010; Jeshion, 2013; Kennedy, 1992; Nunberg, 2018; Popa-Wyatt & Wyatt, 2018).

¹⁶ Slurs: $M_{\text{direct}} 4.01$ vs. $M_{\text{indirect}} 3.37$, $p < .001$ (Cepollaro et al., 2019: 38).

¹⁷ One of the virtues of Prohibitionism is its ability to explain why different embedding cases such as indirect and direct quotation, are still offensive. The claim is that the utterance of the slur is what is offensive because there is a prohibition on such utterances. However, if this is the case, then there is no explanation of why non-embedded cases would be more offensive.

I think this criticism is apt if we think of Prohibitionism as a monistic theory. However, if we take it that there are multiple sources of offense, and the offense can, to some degree, be additive, then we can explain this variation by appealing to these different sources. For example, a prohibitionist account of a slur would have problems explaining the variation of perceived offensiveness between standard and non-standard use cases (in this case, direct and indirect quotation). However, the addition of a doxastic source of offense in standard use cases may increase the offensiveness of a slur utterance, as the study demonstrates.

If I am correct that slur utterances are the locus of a convolution of linguistic mechanisms, a false belief, and shoddy justificatory practices, then a pluralist approach to slurs is both needed and of interest.

Acknowledgements Thanks to Scott Aikin, Ben Lennertz, and Ron Mallon for helpful discussion and feedback.

Funding There were no sources of funding for this research.

Data Availability Not applicable.

Declarations

Ethics Approval Not applicable.

Patient Consent Not applicable.

Permission to Reproduce Material from other Sources No external sources were used.

Conflict of Interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Anderson, L., & Lepore, E. (2013). Slurring words. *Noûs*, 47(1), 25–48. <https://doi.org/10.1111/j.1468-0068.2010.00820.x>
- Ashwell, L. (2016). Gendered Slurs. *Social Theory and Practice*, 42(2), 228–239. <https://doi.org/10.5840/soctheorpract201642213>
- Asshole, n. (2022). In *OED Online*. Retrieved from <https://www.oed.com/view/Entry/47120124#eid121208830>
- Austin, J. (1962). *How To Do Things With Words*. <https://doi.org/10.1093/acprof:oso/9780198245537.001.0001>

- Bach, K., & Harnish, R. (1979). Speech acts. In *Routledge Encyclopedia of Philosophy*. <https://doi.org/10.4324/9780415249126-u043-1>
- Bach, K. (2018). Loaded words: On the semantics and pragmatics of slurs. *Bad Words: Phil Perspect Slurs*, 60–76. <https://doi.org/10.1093/oso/9780198758655.003.0004>
- Blakemore, D. (2015). Slurs and expletives: A case against a general account of expressive meaning. *Language Sciences*, 52, 22–35. <https://doi.org/10.1016/j.langsci.2014.06.018>
- Boisvert, D. R. (2008). Expressive assertivism. *Pac Phil Quart*, 89(2), 169–203. <https://doi.org/10.1111/j.1468-0114.2008.00315.x>
- Bolinger, R. J. (2017). The Pragmatics of Slurs. *Nous*, 51(3), 439–462. <https://doi.org/10.1111/nous.12090>
- Camp, E. (2013). Slurring Perspectives. *Anal Phil*, 54(3), 330–349. <https://doi.org/10.1111/phib.12022>
- Camp, E. (2018). A dual act analysis of slurs. In D. Sosa (Ed.), *Bad words: Philosophical perspectives on slurs*. (Vol. 1). Oxford University Press.
- Cepollaro, B. (2015). In defence of a presuppositional account of slurs. *Language Sciences*, 52, 36–45. <https://doi.org/10.1016/j.langsci.2014.11.004>
- Cepollaro, B., & Stojanovic, I. (2016). Hybrid evaluatives: In defense of a presuppositional account. *Grazer Philosophische Studien*, 93(3), 458–488. <https://doi.org/10.1163/18756735-09303007>
- Cepollaro, B., Sulpizio, S., & Bianchi, C. (2019). How bad is it to report a slur? An empirical investigation. *Journal of Pragmatics*, 146, 32–42. <https://doi.org/10.1016/j.pragma.2019.03.012>
- Crimmins, M. (1992). Tacitness and virtual beliefs. *Mind & Language*, 7(3), 240–263. <https://doi.org/10.1111/j.1468-0017.1992.tb00207.x>
- Croom, A. M. (2011). Slurs. *Language Sciences*, 33(3), 343–358. <https://doi.org/10.1016/j.langsci.2010.11.005>
- Čupković, G. (2015). Diachronic variations of slurs and levels of derogation: On some regional, ethnic and racial slurs in Croatian. *Language Sciences*, 52, 215–230. <https://doi.org/10.1016/j.langsci.2014.09.002>
- DiFranco, R. (2017). Appropriate slurs. *Acta Analytica*, 32(3), 371–384. <https://doi.org/10.1007/s12136-016-0313-0>
- Dummett, M. (1973). Frege and Wittgenstein. In *Frege and Other Philosophers*, pp. 237–248. <https://doi.org/10.1093/019823628x.003.0011>
- Expressives and beyond: An introduction to varieties of use-conditional meaning. (2013). In: Gutzmann, D., Gärtner, H-M. (eds.) *Beyond Expressives: Explorations in Use-Conditional Meaning*. https://doi.org/10.1163/9789004183988_002
- Fletcher, M. A. (1997, October 8). FUROR ERUPTS OVER RACIAL EPITHET. *Washington Post*. Retrieved from <https://www.washingtonpost.com/archive/politics/1997/10/08/furor-erupts-over-racial-epithet/16a5c19b-8c49-42a4-bc70-975e59f990a1/>
- Gray, DM, Lennertz, B. (2020). Linguistic Disobedience. *Philosopher's Imprint*, 20(21), 16. Retrieved from <http://hdl.handle.net/2027/spo.3521354.0020.021>
- Grice, H. P. (1989). *Studies in the way of words*. Harvard University Press.
- Hom, C. (2008). The semantics of racial epithets. *The Journal of Philosophy*, 105(8), 416–440. <https://doi.org/10.5840/jphil2008105834>
- Hom, C. (2010). Pejoratives. *Philosophy Compass*, 5(2), 164–185. <https://doi.org/10.1111/j.1747-9991.2009.00274.x>
- Hornsby, J. (2001). Meaning and uselessness: How to think about derogatory words. *Midwest Studies in Philosophy*, 25(1), 128–141. <https://doi.org/10.1111/1475-4975.00042>
- Bennett Jr., L. (1982). Why Richard Pryor Stopped Saying “Nigger”. *Ebony*, 37(8), 118–120; 122; 124; 126. <https://books.google.com/books?id=cuKlqywEMaQC&printsec>
- Jeshion, R. (2013). Expressivism and the Offensiveness of Slurs. *Philosophical Perspectives*, 27(1), 231–259. <https://doi.org/10.1111/phpe.12027>
- Jeshion, R. (2017). Slur creation, bigotry formation: The power of expressivism. *Phenomenology and Mind*, 11(11), 130–139. https://doi.org/10.13128/PHE_MI-20113
- Karen, n. (2020). In *Dictionary.com*. Retrieved from <https://www.dictionary.com/e/slang/karen/>
- Kennedy, R. (1992). *Nigger: The Strange Career of a Troublesome Word*. Retrieved from <https://www.penguinrandomhouse.com/books/91619/nigger-by-randall-kennedy/>
- Lewis, H. (2020, August 19). The Mythology of Karen. Retrieved 1 March 2023, from The Atlantic website: <https://www.theatlantic.com/international/archive/2020/08/karen-meme-coronavirus/615355/>
- Lycan, W. G. (2015). Slurs and lexical presumption. *Language Sciences*, 52, 3–11. <https://doi.org/10.1016/j.langsci.2015.05.001>

- Mandelbaum, E., Ware, J., & Young, S. (2024). The Sound of Slurs: Bad Sounds for Bad Words. In Joshua Knobe & Shaun Nichols (Eds.), *Oxford Studies in Experimental Philosophy* (1st ed.). Retrieved from <https://global.oup.com/academic/content/series/o/oxford-studies-in-experimental-philosophy-osep/>
- McCready, E. (2010). Varieties of conventional implicature. *Semantics and Pragmatics*, 3(8). <https://doi.org/10.3765/sp.3.8>
- Miller, S. (2019, December 7). My So-Karen Life. *The New York Times*. Retrieved from <https://www.nytimes.com/2019/12/07/style/its-karentown.html>
- Neufeld, E. (2019). Essentialist theory of meaning of slurs. *Philosopher's Imprint*, 19(35). <http://hdl.handle.net/2027/spo.3521354.0019.035>
- Nunberg, G. (2018). The social life of slurs. In D. F. D. Harris & M. Moss (Eds.), *New Work on Speech Acts* (pp. 237–295). <https://doi.org/10.1093/oso/9780198738831.003.0010>
- Popa-Wyatt, M., & Wyatt, J. L. (2018). Slurs, roles and power. *Philosophical Studies*, 175(11), 2879–2906. <https://doi.org/10.1007/s11098-017-0986-2>
- Potts, C. (2007). The expressive dimension. *Theoretical Linguistics*, 33(2). <https://doi.org/10.1515/TL.2007.011>
- Predelli, S. (2010). From the Expressive to the Derogatory: On the Semantic Role for Non-Truth-Conditional Meaning. In Sawyer, S. (Ed.) *New Waves in Philosophy of Language* (164–185). https://doi.org/10.1057/9780230248588_9
- Richard, M. (2008). *When Truth Gives Out*. <https://doi.org/10.1093/acprof:oso/9780199239955.001.0001>
- Saka, P. (2007). *How to Think About Meaning*. <https://doi.org/10.1007/1-4020-5857-8>
- Schlenker, P. (2007). Expressive presuppositions. *Theoretical Linguistics*, 33(2). <https://doi.org/10.1515/TL.2007.017>
- Sennet, A., & Copp, D. (2015). What kind of a mistake is it to use a slur? *Philosophical Studies*, 172(4), 1079–1104. <https://doi.org/10.1007/s11098-014-0338-4>
- Spotorno, N., & Bianchi, C. (2015). A plea for an experimental approach on slurs. *Language Sciences*, 52, 241–250. <https://doi.org/10.1016/j.langsci.2015.04.004>
- Thief, n. (2022). In *OED Online*. Retrieved from <https://www.oed.com/view/Entry/200740>
- Tirrell, L. (1999). Derogatory Terms: Racism, Sexism and the Inferential Role Theory of Meaning. In K. Oliver & C. Hendricks (Eds.), *Language and Liberation: Feminism, Philosophy and Language* (pp. 41–79). SUNY Press.
- Truedell, K. (Director). (1996). *Chris Rock: Bring the Pain* [Comedy Special]. Retrieved from <https://www.dailymotion.com/video/x65kk41>
- Whiting, D. (2013). It's Not What You Said, It's the Way You Said It: Slurs and Conventional Implicatures. *Analytic Philosophy*, 54(3), 364–377. <https://doi.org/10.1111/phib.12024>
- Williamson, T. (2009). Reference, Inference, and the Semantics of Pejoratives. In Joseph Almog & Paolo Leonardi (Eds.) *The Philosophy of David Kaplan*, pp.137–158). <https://doi.org/10.1093/acprof:oso/9780195367881.003.0009>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.