

Slurs, Inflammatory Language, and the Specificity Problem

ROBIN JESHION

University of Southern California, Los Angeles, USA

In Inflammatory Language, Una Stojnić and Ernie Lepore argue that no extant theory of slurs can explain slurs' hyperprojectivity, emphasizing their difficulties in accounting for acoustic and phonological resemblance cases in which a word merely sounds like a slur. Further, all content theories confront the Specificity Problem, the charge that the content view's content, whatever it is, is too specific to encompass the full range of competent weapon uses of slurs. One half of this paper concerns hyperprojectivity. I argue that there is a gap in Inflammatory Language's overarching dialectic that results from excluding a range of theories. Some theories of slurs are what I call single mechanism views: they aim to explain all the phenomena with a single explanatory mechanism. Multiple mechanism views exploit more than one. Within Inflammatory Language, multiple mechanism theories are bypassed. Yet multiple mechanism theories possess resources to explain slurs' hyperprojectivity. The other half of this paper addresses the Specificity Problem. I argue that a view I have developed in previous writings, Identity Expressivism, does not succumb to the problem. I craft a version of the Specificity Problem tailor-made for the theory and rooted in Stojnić and Lepore case against other expressivist theories. Identity Expressivism is, I argue, uncompromised by the Specificity Problem.

Keywords: Slurs; semantics; pejoratives; epithets; expressives; hate speech.

1 The overarching dialectic of Inflammatory Language

Within their rich investigation of slurs in *Inflammatory Language*, Una Stojnić and Ernie Lepore advance both a negative and a positive thesis. Their negative thesis is that no extant theory of slurs can explain slurs' hyperprojectivity. Their positive thesis is that the Articulation Account

they develop within *Inflammatory Language* is capable of doing so.¹ This paper dominantly concerns their negative thesis. Here, I argue that their negative thesis does not go through.

I focus on Stojnić and Lepore's arguments for the negative thesis as applied to content theories. They advance two main arguments. The first is that all content theories are unable to explain slurs' hyperprojectivity, emphasizing inheritance cases on which a word sounds like a slur. The second is that all content theories confront the *Specificity Problem*, the charge that the content view's content, whatever it is, is too *specific* to encompass the full range of competent weapon uses of slurs.

The first half of this paper addresses the claim that content views are bereft of resources to explain hyperprojectivity. I argue that there is a gap in *Inflammatory Language's* overarching dialectic that results from excluding a range of theories. Some theories of slurs are what I call single mechanism views: they aim to explain all the phenomena with a single explanatory mechanism. Multiple mechanism views exploit more than one. Within *Inflammatory Language*, multiple mechanism theories are bypassed. The reason why appears to be that Stojnić and Lepore either assume that all theories must be single mechanism or illegitimately rule out of hand multiple mechanism approaches. Yet multiple mechanism theories possess resources to explain slurs' hyperprojectivity.

The second half of this paper concerns Stojnić and Lepore's Specificity Problem. My aim is to demonstrate that the view I have developed in previous writings, Identity Expressivism, does not succumb to the problem. While Stojnić and Lepore don't directly address Identity Expressivism, I craft a version of the Specificity Problem tailor-made for the theory and rooted in Stojnić and Lepore case against other expressivist theories. Identity Expressivism is, I argue, uncompromised by the Specificity Problem.

2.1 Slurs' hyperprojectivity and mere orthographic and acoustic resemblance cases

Stojnić and Lepore frame their investigation in *Inflammatory Language* as one that primarily aims to explain what they call slurs' 'pejorative sting'. They ask: "What is the nature and source of [slurs'] pejorative sting, that makes slurs such powerful linguistic weapons?" (Stojnić and Lepore 2025: 2). In keeping with much of the slur's literature, they aim to account for how slurs function as tools of derogation, their capacity to cause offense when used literally as weapons, their derogatory and

¹ In Jeshion (2025), I discuss the positive thesis. There I explicate why the Articulation Account cannot explain how slurs function as tools of derogation, cannot account for instances of incomplete understanding of slurs, and is hard-pressed to analyze slurs' linguistic standing as pejorative expressions. I also raise concerns about some of their data.

offense potential when used within attitude attributions, in direct and indirect quotation, and even when the slur is mentioned. As noted, they advance two theses, one negative, one positive. The negative thesis is that no extant view in the contemporary literature is adequate to the task of explaining the full range of phenomena of how slurs sting. They break these down into two types, what they call content views and non-content views. Content views construe slurs as encoding or conveying a pejorative meaning or message. Different content views assign a different variety of content. Semantic content theories invoke a description, expressive, or perspective as a slur's meaning or within its use conditions. Pragmatic content theories typically convey descriptive, expressive, or ideological affiliation, yet do so via pragmatic mechanisms, like presuppositions or generalized conversational implicatures.² All content theories, say Stojnić and Lepore, are defective because "there is no pejorative content that can capture the behavior of a slur's offensive potential" (Stojnić and Lepore 2025: 2).

Non-content views make no appeal to meaning or pragmatically conveying a content. They appeal, rather, to sociolinguistic features of slurs or to the psychological impact of hearing a slur to account for their 'offensive potential'. For instance, Prohibitionism locates slurs' sting in their taboo status, in the offensiveness of breaking prohibitions on using or saying slurs (Anderson and Lepore 2013). Word Associationism locates that sting in associations triggered upon seeing or hearing a slur (Lepore and Stone 2014). For both non-content views, *slurring words themselves* are the primary source of slurs' sting. For Prohibitionism, speakers offend by uttering prohibited words. For Word Associationism, speakers trigger associations in their audience by uttering slurs. To Stojnić and Lepore, therein lies a problem common to all extant non-content theories, the fact that they ultimately root slurs' pejorative sting in slurring word themselves.

Why is there no pejorative content that can explain slurs' offense potential? And what is the error in rooting slurs' sting in slurring words themselves? According to Stojnić and Lepore, content and non-content views confront cases of pejorative sting via the *mere orthographic or acoustic resemblance to a slur*. These are their inheritance cases. Here is one: in 1999, a white aide to the mayor of Washington DC used the expression *niggardly* to describe a budget. Although the word shares orthographic and acoustic resemblance with the racial slur, it is unrelated etymologically and semantically. Nevertheless, upon hearing the aide's utterance, one of his black colleagues took offense, lodging a complaint. A similar case occurred in 2020, at my own university, USC. While conducting a lecture about cross-cultural dialogue, a communications professor noted how different languages employ different filler

² Some content views: Semantic: Bach (2018), Camp (2013), (2018), Davis and McCready (2020), Hom (2008), Hom and May (2013), (2015), Jeshion (2013b), (2018), Marquez and Garcia-Carpintero (2020), Potts (2007), Richard (2008). Pragmatic: Jorgensen Bolinger (2017), Nunberg (2018).

expressions. In English, we use ‘um’ and ‘er’. Speakers of Mandarin, he noted, use a filler term, ‘那个’, an expression that is pronounced like the n-word. In the lecture, the professor enunciated the word. He did so with no intention to provoke – indeed, apparently entirely oblivious to its possible effects. A group of students in the course expressed outrage over the incident, saying it caused them pain and upset. Call these *mere orthographic/acoustic resemblance cases*. Call the problem to explain all the ‘offense potential’, where this includes explaining derogation of weapon uses, derogatory variation, offense of belief attributions, mere mentions, and so on *the problem of slurs’ hyperprojectivity*.

Of mere orthographic/acoustic resemblance cases, Stojnić and Lepore write:

Such cases are particularly puzzling for content-based accounts: the offensive potential is inherited even though no slur is either used or mentioned. Indeed,...this type of data presents an issue for virtually all extant theories – content or non-content. The problem is that the expressions that inherit the pejorative effect are distinct from slurs and do not share either content or etymology with them. So, any account that ties the pejorative effect to an expression will leave unexplained why merely resembling an articulation of a slur leads to a pejorative effect being inherited by tokenings of a resembling expression.... (Stojnić and Lepore 2025: 75)

They add that for content views specifically, orthographic/acoustic resemblance cases present an additional issue: “the explanation of inheritance cannot be gotten by way of meaning – for the expressions that inherit the pejorative effect don’t share any aspect of the slur’s meaning” (Stojnić and Lepore 2025: 75).

Stojnić and Lepore’s other thesis is positive: that the novel view they craft and dub the Articulation Account can explain all the phenomena. According to the Articulation Account, the offensive potential of slurs resides in negative association triggered by

certain of [a slur’s] articulations – its phonological and/or orthographic forms. These associations attach to particular articulations of slurs through multifarious factors – causal, historical, cultural, and psychological. Token articulations trigger them even when they do not accompany a tokening of a slur, while tokenings of slurs, absent particular standard articulations that harbor these associations, fail to trigger them. In short, the offensive potential has nothing to do with slurring words; it is a result of associations triggered by articulations of slurs. (Stojnić and Lepore 2025: 2)

For them, articulations and not slurring words are the ultimate source of slurs’ sting. The theory is, then, expressly designed to account for mere orthographic/acoustic resemblance cases. Since, in the case of the Mandarin filler term it is, they say, the sheer acoustic resemblance that provoked the outrage, they invoke the sound itself as what accounts for slurs’ ‘offense potential’. They maintain further, that by appealing *exclusively* to articulations, the Articulation Account successfully accounts for all the phenomena, including all aspects of slurs’ hyperprojectivity. This is their *Comprehensiveness Thesis*. They also maintain that their positive account is “uniquely well-positioned to account for

the full range of data and do so in a fully uniform manner” (Stojnić and Lepore 2025: 6). This is their *Uniformity Thesis*.

2.2 *Single mechanism and multiple mechanism theories of slurs*

Implicit in Stojnić and Lepore’s *Comprehensiveness* and *Uniformity* Theses is a commitment to what I call a single mechanism approach to theorizing about slurs. As its name suggests, single mechanism accounts posit a single mechanism to explain all the phenomena regarding the full range of slurs’ hyperprojectivity. The single mechanism that Stojnić and Lepore embrace to explain slurs’ pejorative potential is articulations. Other theorists also appear to plumb for a single mechanism approach. For instance, Anderson and Lepore emphasize but one mechanism – prohibition-breaking – to explain hyperprojectivity (Anderson and Lepore 2013).

By contrast, multiple mechanism approaches posit various mechanisms to explain the full range of phenomena regarding the derogatory and pejorative potential of slurs. In earlier papers, I advanced a multiple mechanism theory of slurs (Jeshion 2013b, 2018). The theory advocates an Identity Expressivist semantics of slurs whose primary role is to capture why slurs are, by convention, pejorative expressions, and how, in weapon uses speakers derogate their targets, a matter fully independent of whether anyone takes offense to the act or any downstream causal harmful and painful effects due to the utterance (Jeshion 2013b). The account couples the expressivist mechanisms in the semantics with additional pragmatic and sociolinguistic features of slurs, for there are a “plethora of reasons why a use of a slurring term can be offensive, and we must tease apart the sources of offensiveness of slurring terms due to their semantic properties and those that are due to pragmatic phenomena and sociolinguistic properties” (Jeshion 2013b: 234).

I advanced several mechanisms contributing to the offense profile of slur-utterances that supplement those that derive exclusively from the semantics. One was stereotype-activation. While, I argued, slurs do not semantically encode stereotypes, they have the power to activate, to trigger, thoughts of stereotypes in hearers, and cause harm and offense due to that activation (Jeshion 2013a, 2013b). Another mechanism concerned how utterances of slurs interact with historical institutional bigotry and oppression of targets: in general, utterances of slurs will compound the effects of oppressive ideologies and will cause far more psychological and social damage to those in historically oppressed groups. Still, because the mechanism is causal and associationist, the exact impact will vary radically from utterance to utterance, and from slur to slur. In addition to stereotype activation and interaction with ideologies and historical oppression, other psychological and social factors can contribute to intensifying or diminishing the impact of a slur.

The main message: as always with perlocutionary effects, downstream discourse effects are wide-ranging, manifold, complicated. And finally, I argued that Lepore and Anderson were right to identify prohibition-breaking as important. While it is just one piece of the story and cannot explain derogation, breaking the taboo of uttering a slur does cause offense to those invested in the prohibition's integrity. Thus, another sociolinguistic mechanism, prohibition-breaking, contributes to the offense profile of slur utterances (Jeshion 2013b). Several other theorists, including Diaz-Legaspe, Liu, and Stainton, Rappaport, and Rinner and Hiecke have also advanced their own multiple mechanism theories or at least advocated handling some of the hyperprojective phenomena with multiple mechanisms.³

2.3 *Exclusion of multiple mechanism accounts*

Return now to *Inflammatory Language's* negative thesis that no extant theory can explain slurs' hyperprojectivity. The most important thing to recognize about Stojnić and Lepore's dialectic is that all of arguments are directed at single mechanism views. They curiously ignore multiple mechanism theories. Moreover, throughout the book, they repeatedly deploy a mode of argumentation that appears to rule out of hand multiple mechanism approaches – perhaps for an underlying rationale that they do not specify. Alternatively, they may simply implicitly assume that all theories must be single mechanism. In fact, in their critical discussion of many different views, they never allow for a theory to tack on an additional mechanism to explain aspects of hyperprojectivity, including mere orthographic/acoustic resemblance cases. This is precisely what Rappaport does to account for the toxicity in mere mentions, positing that slurs undergo distinct neurolinguistic processing, routed through their phonological forms. He could easily deploy it as well to explain mere acoustic resemblance cases. The move seems in no way implausible, and certainly not illegitimate *for* simply adding an additional explanatory mechanism.

Let's capture Stojnić and Lepore's implicit assumption thus:

Single Mechanism Necessity Assumption: if a theory is unable to explain hyperprojectivity with a single mechanism, that theory cannot explain hyperprojectivity.

³ Diaz-Legaspe, Liu, and Stainton (2019) offer a register-based account of the conventional rules governing slurs, which would count as a content view. It is multiple mechanism because they acknowledge that nonlinguistic performance effects are necessary to fully account for hyperprojectivity. Rappaport (2019) advances a theory that couples together a Camp and Nunberg style group-allegiance or perspective signaling plus a subsidiary neurological processing mechanism to explain the toxicity of slurs in indirect quotation and slur mentions. Rinner and Hiecke (2021) do not themselves advance a particular semantic or pragmatic analysis of slurs. Nevertheless, they argue that content views can handle derogation, while another mechanism handles slur mentions, namely the psychological efficacy (via associations, presumably) of being reminded of the content.

The Single Mechanism Necessity Assumption is false. After all, there are several multiple mechanism theories of slurs that, in addition to positing a content, advance other mechanisms – hearers’ associations, neurolinguistic factors, stereotype activation, historical facts, and sociolinguistic properties of words – to explain hyperprojectivity. They are not, *for that*, unable to explain hyperprojectivity. And they are not, *for that*, inadequate theories.

So, the Single Mechanism Necessity Assumption is false. For this reason, one might naturally be skeptical of the claim that numerous arguments in *Inflammatory Language* depend upon it. To back this claim, then, consider the following three arguments.

First, addressing Potts’ expressivist conventional implicature theory,⁴ Stojnić and Lepore note that different slurs for the same group vary in their pejorative effects and impact. They write:

This is particularly difficult to capture on attitudinal or fully non-descriptive approaches. Such accounts *have to argue that distinct slurs for the same target group signal different negative attitudes of differing strength*, in way that underscores speaker intuitions about the associated pejorative effect. We find it highly dubious that fine-grained distinctions in attitudes or affects ...underwrite such intuitions. We are thus pessimistic about the success of CI accounts. (Stojnić and Lepore 2025: 51, my emphasis)

This argument presumes that Potts has no resources available to him other than his semantics to explain derogatory variation of different slurring words for the same group.

Second, addressing Camp’s perspectivalist theory,⁵ Stojnić and Lepore note that slurs often “evoke facts that go well beyond” that which is encoded in perspectives, in particular they may invoke “socio-historical facts about who used or coined the term, including socio-cultural facts concerning power structures and structures of oppression; or various phonological/orthographic features”; or they may evoke imagery. The only way for Camp to rectify her theory, they maintain, is to widen perspectives to incorporate “all open-ended associations that can affect pejorative potential.” This, they say, only exacerbates problems by bloating the semantics, making disagreements in pejorative effect linguist disagreements (Stojnić and Lepore 2025: 59–61).

Lastly, in summarizing their case against content views, Stojnić and Lepore claim that all such views are faced with an insurmountable problem of slur-meaning insulation, namely, that in quotation and in mere displays of slurs, the slur’s meaning is inert. They write:

we have seen that the projective behavior of the pejorative potential is way more robust than any of those accounts – semantic or pragmatic – predict. Indeed, the effect remains even in environments that render meaning inert – meaning attributions, quotative environments, and even mere displays of slurs. In as much as these environments render meaning inert, that they

⁴ Potts (2005), (2007).

⁵ Camp (2013), (2018).

still preserve the offensive potential of slurs suggests that the latter *cannot be cashed out in terms of meaning*. (Stojnić and Lepore 2025: 129)⁶

Their concluding sentence, that the persistent sting of slurs within meaning-insulated constructions *cannot* be explained in terms of meaning is presented here as a nail in the coffin of all content theories.

In each of the three arguments, Stojnić and Lepore *assume* that the only mechanism available to explain derogatory variation, offense of a use or mention, or same-sounding articulation is the semantic mechanism the theorist embraces: expressivist conventional implicature, for Potts; perspectives for Camp; and for all other content theories, whatever mechanism is exploited to capture content. From this, they conclude – rightly, I maintain – that content, however it is cashed out, is insufficient to account for *all* the phenomena. Yet they then conclude from *that*, that the theory itself is thereby defective. This last inference depends essentially on the false *Single Mechanism Necessity Assumption*.

Now, what general lessons we can learn from this dialectical gap and assumption within *Inflammatory Language*? I offer three. First, *assuming* that the goal of our communal project of understanding slurs is to capture all the effects, intended and unintended, of every communicative act and every sound people make that involves the assertion or mention of a slur and of all intentional and unintentional acts involving humans making noises that sound like a slur, then every theory can – and I think should – reject a single mechanism approach. No content theorist is *restricted* from supplementing their account by positing pragmatic and sociolinguistic mechanisms to explain this extraordinarily vast range of discourse effects. Further, it should not be assumed that the mechanisms used to explain how slurs derogate and cause offense be exactly the same as those used to explain why homonyms can trigger offense.

Second, this raises an interesting question about the proper methodological principles for multiple mechanism theories. Obviously, appealing to one mechanism is always acceptable. Equally obviously, adding mechanisms *willy-nilly* is unacceptable. Advancing and justifying guiding principles and rationales is a large task, not one I can take up fully here, though some of §3.1 speaks to this matter.

Last, we must take stock and ask whether our communal project ought to be this vast in its scope. Is it in fact incumbent upon a theory of slurs to account for all the intended and unintended discourse effects of utterances and mentions of slurs and all articulations that may sound like slurs? While this is an important methodological question, here is not the place to take a deep dive into what, precisely, our com-

⁶ This style of argument also occurs in numerous other places throughout the book. Cf., chapter three, where the authors maintain that the only way for a presuppositional theorist to account for the different discourse effects of different slurs is to “posit different pejorative presuppositional contents for different slurs.”(42) Cf., also their footnote 75, (61–62).

munal project should be. But I will register my skepticism that it is incumbent upon us to grapple with every effect of a speech act that looks or sounds like a slur and, in particular, to treat it as on a par with intentional weapon uses of such terms. What is puzzling is that mere orthographic/acoustic resemblance cases are the prime variety of datum inspiring Stojnić and Lepore's Articulation Account. In my view, this is misguided. As I noted elsewhere, outrage in response to slur homophones that are not the n-word are exceptionally rare.⁷ Further, it seems to me a banal fact about communicative acts, and intentional acts more generally, that those that appear to listeners to have the same properties of those that systematically cause offense will also sometimes trigger that same effect. Such phenomena shouldn't induce novel theory.

In sum: Stojnić and Lepore's negative thesis is that all extant theories are unable to explain slurs' hyperprojectivity. Content views cannot appeal to content to explain why utterances on non-slurring words that sound or look like slurs trigger offense. And since both content and non-content views maintain slurs themselves – the words, not their articulations – are the ultimate source of their pejorative sting, both cannot explain the offense triggered in mere orthographic/acoustic resemblance cases. Thus far, I have taken a preliminary step toward undermining this negative thesis by arguing that Stojnić and Lepore suppose that all theories must deploy but a single explanatory mechanism. This assumption is unargued and false. Any particular multiple mechanism theory of slurs may well be wrong. But they are not wrong for positing multiple mechanisms. This leaves a path open for the viability of multiple mechanism theories.

3 The Specificity Problem: A defense for Identity Expressivism

In this second section, I extend the case against the negative thesis by demonstrating the power of my favored multiple mechanism theory, a view I call Identity Expressivism.⁸ On Identity Expressivism, slurs are tools of derogation. In weapon uses of slurs, speakers regard their targets with contempt, and with the slur express their contempt. It counts squarely as a content view. Stojnić and Lepore's most powerful arguments against all content views is given by their *Specificity Problem*. This is the charge that the content in any given content view is too *specific* to encompass the full range of competent weapon uses of slurs. I demonstrate that key features of Identity Expressivism disable the Specificity Problem. I begin by briefly laying out four methodological assumptions I adopt. Next, I give a bare-bones overview of Identity Expressivism, and then turn to the Specificity Problem.

⁷ Jeshion (2025).

⁸ The view is developed in Jeshion (2013b), (2016), (2017), (2018).

3.1 Methodological assumptions of Identity Expressivism

Different methodologies, starting points, and underlying assumptions often strongly impact an analysis of slurs. Although I cannot detail the whole range or offer full justifications for them, here I lay out four of my own, all of which I believe are or ought to be shared by Stojnić and Lepore. Each plays a critical role in the dialectic regarding the Specificity Problem.

One: Slurs are pejorative lexical items.

From a linguistic perspective, there are two basic starting points for examining slurs. One is to view slurs, first and foremost, as *types of lexical items*. Another is to take as most basic a type of speech act, the act of slurring, where such acts can occur in the absence of a word standardly branded as a slur. I adopt the first perspective. This approach affords insight into how slurs are related to other lexical items, including other pejoratives and dysphemisms, and also kind terms, names, and the many types of social deictics, including honorifics, pronouns indicating (in)formality, nicknames, diminutives, and other hypocorisms. Critically, this starting point allows us to understand slurs' lexical dynamics: how they can change their linguistic properties over time and become polysemous.⁹

Two: Basic uses of slurs are weapon uses, applied to the target group.

Slurs are used in various ways. Notoriously, slurs are used to derogate and dehumanize. Yet they are also used, typically by target group members themselves, to neutrally reference the group or group-member or to convey positive valences of solidarity, endearment, friendship, or empowerment. Slurs are also used to express bigotry that is exception-making and to widen the domain of application beyond the target group.

Ideally, a full theory of slurs ought to have the resources to explain *all* of the phenomena. However, no linguistic theory that isn't hyper-contextualist will be able to satisfy all the phenomena without appealing, at some point, to slurs' polysemy. I adopt the following approach: in basic uses, slurs are used to pejoratively reference an important social group, defined by race, ethnicity, nationality, religion, gender, sexual orientation, and so on. Often such social groups have a corresponding non-derogating term, a neutral counterpart, as a preferred term for the group – but not always.

Basic uses have two properties. They function to derogate. They are thus aptly deemed *weapon uses*. Basic uses also aim to refer to all and only those in the group that the slur targets. I call such uses *group-*

⁹ Stojnić and Lepore treat slurring terms as first and foremost lexical items, more primitive than slurring speech acts. They announce that 'slurs are pejorative by design'(2). However, it is unclear how, precisely, they explain slurs' standing as *pejorative* lexical items. Cf., Jeshion (2025).

referencing (G-referencing) uses of slurs. Identity Expressivism's semantics is designed to apply to these basic uses.

While I construe group-referencing weapon uses as basic, Identity Expressivism is also designed to account for *non-basic uses*, like those with positive and neutral valence and with different extensions (what I call G-contracting and G-extending uses). I explain non-basic uses as securing what they convey via novel extensions from basic uses. The theory uses familiar tools of semantic change and polysemy to explain this multiplicity of meanings.

Note: our methodological assumption that selects a class as basic does *not* entail that weapon group-referencing uses are the only uses whose meanings are conventionalized. The n-word has a conventionalized meaning as a social deictic to convey friendship. Similarly, our methodological assumption does not entail that for every slurring word, its weapon group-referencing use is currently or at some other time the most common use of that expression. For instance, the slurring word *queer* has a basic use, a weapon group referencing use, whereby it functions as a linguistic tool to derogate LGBTQ+ persons. Through most of the twentieth century, in most western communities, the basic use was the dominant use. With the reclamation of *queer* beginning in the late 1980s, the word became polysemous: the neutral use designating a group with a novel identity – *being queer* – proliferated, becoming conventionalized. Currently, at least in many western locales one quarter way through the twenty first century, the basic use is, happily, no longer the dominant use. This is all compatible with Identity Expressivism.

Three: Sharply distinguish derogation, offensiveness, and the moral assessment of uses of a slur.

I sharply distinguish between *derogation* and *offense*.¹⁰ I use *derogate* and cognates to label what speakers do – what they perform – when deploying slurs as weapons. Such speakers derogate their targets and they do so regardless of how anyone responds to the act. Derogation is, therefore, not the same as *offense taken*, understood as an actual response to an utterance. It also differs from *offense potential*, understood as a possible response to an utterance or as a measure of possible responses. Offense taken and offense potential both concern *post-act* (actual and potential) *responses* to a slur content, utterance, mention, attribution, or even to non-slur articulations, utterances that just sound or look like a slur.

¹⁰ Officially, Stojnić and Lepore would appear to agree. Slurs, they say, are “epithets that *derogate* purely on the basis of group membership, e.g., on the basis of race, ethnicity, origin, religion, gender, sexual orientation, or ideology.” (2, emphasis mine) and are “used to derogate, disparage, offend, insult, or cause harm.” (5) Nevertheless, they standardly conflate the distinction between derogation and offense with a covering term ‘pejorative potential’ or ‘offensive potential’. This creates problems for their Articulation Account’s ability to explain cases (below) like *Homophobia*, a point argued in Jeshion (2025).

I sharply separate the *moral assessment* of an act of using a slur from the non-normative account of what speakers do with slurs. That is, to say that a speaker derogates or offends with a slur is not *thereby* to say that their act is worthy of moral criticism or condemnation. That is a further matter.

Four: A linguistic theory of slurs must be capable of explaining both derogation and offense.

All agree that slurs are tools of derogation. In the vast range of basic uses of slurs, speakers derogate *and* cause offense. How can we separate the two? The best way is to consider cases where a speaker derogates with a slur while the actual and potential offense are made null. Here is one:

Homophobia: A and C are friends, both homophobic. Their homophobic attitudes are mutual knowledge between them. They frequently use *f****t* amongst themselves to designate people they deem gay. Both A and C believe that B is gay. Both know that B does not understand the slur – for, they know, B barely understands English. In isolation of all other people except C, A says to B, ‘You are a *f****t*’. Since B has never heard the slur before, he has no associations whatsoever with it. Going forward, B will never be in contact with anyone who does have any associations. A and C know this fact.

Given the set-up, there is neither any actual offense taken nor any offense potential. Yet with the slur, A has derogated B. This is palpable and should be uncontroversial. That there is derogation here is rooted straightaway in intuition. It employs no theory-laden assumptions about the mechanisms to account for the derogation. Theories of slurs ought to be able to explain the derogation in cases like *Homophobia*.

3.2 Identity Expressivism

Identity Expressivism advances a semantics for what I call canonical slurring terms, those attacking race, nationality, religion, sexual orientation, ability, political affiliation, and possibly other important social categories.¹¹ In their basic uses as weapons, slurs have three components to their semantics:

Group Designating Component: A slurring term designates the group G or set of individuals that is designated by its neutral counterpart. Slurs are designationally-equivalent with their neutral counterparts. The group-designating component contributes to the

¹¹ I contrast canonical slurs from descriptive slurs like *beaner* for Hispanic and *slanty-eyes* for Asians. Like canonical slurs, descriptive slur refer to a target group, yet they also contain an additional stereotype that is truth-conditionally inert. *Beaner* can be ‘aptly’ applied to Hispanic persons that do not eat beans. I also contrast canonical slurs with gendered slurs like *bitch*, *c**t*, *slut*, for girls and women, and *sissy* and *pussy* for boys and men.

truth-value of propositions in the same way that neutral counterparts do.

Expressive Component: Slurring terms are vehicles for expressing affective stances and attitudes: with a slurring term, speakers express contemptuous regard for the members of a group G on account of their being in G or having a group-defining property g. They express that the target is lesser, on a moral dimension, on account of being in the group. The expressive component makes no contribution to truth-conditions.

Identifying Component: Slurring terms are tools for specifying, what, by the speaker's lights, the targets of the slur *are*. Speakers use slurs to map a negative characteristic-defining *social identity* onto the members of the group G. The identifying component makes no contribution to truth-conditions.

For example, [1] and [2] are designationally and truth-conditionally equivalent,

[1] Jake is a Kike

[2] Jake is a Jew

Yet *kike* differs from *Jew* insofar as it is used to express contempt for Jews on account of being Jews. In expressing contemptuous regard, the speaker of [1] does not *say* or *assert* the truth-evaluable proposition

[3] Jake is a person deserving of contempt on account of being Jewish.

And neither does the speaker *say* or *assert*

[4] Jake has the fundamental negative character-defining feature of being Jewish.

It is rather that *by* expressing an affective stance of contempt for Jews on account of being Jews, the speaker displays that what the speaker is, *qua* person, is a Jew and thereby lesser.¹² In this way, the semantics of [1] is essentially equivalent to that of [5] and [6],

[5] Jake is a goddamn Jew

[6] Jake is a Jew^C

where *goddamn* in [5] functions as a pejorative expressive modifier, and the superscript 'C' in [6] indicates that *Jew* is spoken with contemptuous intonation.

Lastly, in giving this expressivist semantics of slurs, we have yet to say how slurs derogate. After all, words – particular lexical items – are not by themselves derogations. *Derogations are acts*, communicative acts, performed by people. Identity expressivism accounts for how, with a basic use of a slur, speakers derogate their targets. In sincerely using a slur with full understanding, a speaker expresses contempt for the

¹² These characterizations leave off many fine-points. Cf., Jeshion (2013b) and (2018) for far more depth.

target, and consequently both *communicates* and *treats* the target as a lesser person, as less deserving of full respect, on account of being in the group.¹³

3.3 *The Specificity Problem*

Finally, we return to the Specificity Problem. Recall that Stojnić and Lepore advance it as the challenge that the content view's content, whatever it is, is too *specific* to encompass the vast range of competent uses of slurs. Importantly, unlike their point about slurs' hyperprojectivity, the Specificity Problem is directed toward basic weapon uses of slurs: not reclaimed uses, not belief attributions, not mentions, not words that sound like slurs. To them, the Specificity Problem undermines all content views, whether the content is descriptive, a speaker's attitude, an emotion, or an ideology encoded as a perspective. Since Stojnić and Lepore don't themselves address Identity Expressivism, I detail how they wield the Specificity Problem against semantic accounts that also posit a close connection between slurs and the emotions: Schlenker's attitudinal presuppositional account of slurs and versions of expressivism due to McCready and Potts. From this, I craft the strongest version of the Specificity Problem for Identity Expressivism, and demonstrate why it fails to undermine the theory.

In *Inflammatory Language*, Stojnić and Lepore present a version of the Specificity Problem to a view that bears *some* similarities to Identity Expressivism, Schlenker's attitudinal presuppositional account of slurs.¹⁴ On Schlenker's view, a slur is truth conditionally equivalent with its neutral counterpart. Here it aligns with Identity Expressivism. Schlenker's analysis differs from Identity Expressivism along several dimensions. For him, the slur is governed by a presupposition that the speaker believes that members of the group the slur targets are despicable. Thus, while Identity Expressivism offers a rule of use to characterize the semantics, Schlenker posits a presupposition as the mechanism to convey the speaker's attitude. More important with regard to the Specificity Problem, Schlenker's view requires the speaker has a standing belief that members of the target group are despicable. Identity Expressivism makes no such assumption. It trades on speakers having affective stances on the group, not any particular belief about them. Finally, Schlenker's belief content is indeed quite specific: that the target group and its members are despicable.

In pressing the problem to Schlenker's account, Stojnić and Lepore claim that it is coherent and perfectly felicitous to use the slur while issuing a denial of the alleged presupposition (Stojnić and Lepore 39-43). Here are two of their key examples against Schlenker:¹⁵

¹³ For a much fuller treatment of the moral psychology and how dehumanization is performed, see Jeshion (2018).

¹⁴ Schlenker (2007).

¹⁵ Here and throughout the rest of this section, my examples are extracted directly from *Inflammatory Language*, yet are altered by shifting the slur, and the

- [7] Jake is a kike, but I don't believe that kikes are despicable.
 [8] There are so many f****ts in that bar. But I don't believe that f****ts are despicable.

Though my reasons for finding [7] and [8] coherent may well be different from theirs, I agree with Stojnić and Lepore. In each utterance, the speaker can be competently using the slur as a weapon while denying they believe that Jewish or gay people are despicable. *Believing* that a whole group is despicable is one thing. Having an *affective stance* of contempt toward a group is another. Contemptuous regard may exist in the absence of standing beliefs that a group is despicable – or, in fact, even that the group is contemptible. And thus, using a slur is compatible with outright sincere denials that the target group is despicable. In my view, this version of the Specificity Problem appropriately pinpoints my own prime concern with Schlenker's account, namely that competent use of a slur does not require that speakers have any specific negative *beliefs* about the group – and further, that such beliefs are not, in the first instance, what is conveyed or otherwise communicated with slurs.

Stojnić and Lepore also wield the Specificity Problem towards views considerably closer to Identity Expressivism. One is McCready's expressivist view, whereby speakers using slurs express negative attitudes toward the slur's target. The problem is that “predications of slurs in the absence of such attitudes are not linguistically infelicitous”.¹⁶ To demonstrate competent use of a slur in the absence of negative attitudes, they offer examples like these:

- [9] I don't think C***ks are despicable. [49]
 [10] John is a n****r. N****rs are very respectable people. [49]
 [11] F****ts are good! We should respect them. [49]

Of [9]-[11], Stojnić and Lepore say they are neither linguistically incoherent, linguistically deviant, nor linguistically infelicitous.

The other view they address is inspired by Potts' treatment of expressives (Potts 2005, 2007) like *damn* when used as intensifiers to modify nouns, as in

- [12] That damn dog is barking again!

According to Potts, the speaker must be “undergoing a heightened, negative affect towards the dog”. On such a view of slurs, in

- [1] Jake is a kike

the speaker expresses a heightened negative affect toward Jews. To such a view, Stojnić and Lepore write

person. They are also trimmed to excise redundancy. I choose to insert slurs in the example sentences, replacing Stojnić and Lepore's Harry Potter 'mudblood', for several reasons. Chief among them is that I had to myself reread the sentences with real slurs in order to test them for felicity. Mistakes in felicity judgments can far more easily arise from considering fictional rather than real-life slurs.

¹⁶ Cf., McCready (2010), Stojnić and Lepore (49).

one needn't have a negative affect toward the target group – not even a standing one, let alone occurrent – in order to felicitously utter the term. We can imagine a casual conversation among bigots, who invariably drop slur terms to reference the target group, regardless of whether they are currently experiencing any negative affect....Such utterances are perfectly felicitous when conjoined with denials of negative attitudes and affects. This, however, is not so for ordinary expressives, such as 'damn', which are infelicitous when conjoined with any such denial. (Stojnić and Lepore 2025: 50).

Here are the contrast cases Stojnić and Lepore advance to support their objection:

[13] Leo is a Kike. But, to be clear, I have no ill feelings toward kikes.

[13'] #Leo is a goddamn Jew. But to be clear, I have no ill feelings toward Jews.

[14] You are a f****t. I love f****ts, they are wonderful people!

[14'] #You are a goddamn homosexual. I love homosexuals, they are wonderful people! (Stojnić and Lepore 2025: 50, 76).

According to Stojnić and Lepore, the predications with the expressive intensifier *goddamn* modifying *Jew* and *homosexual* become linguistically infelicitous with avowals of respect and love, as in [13'] and [14']. Here they agree with Potts. However, they maintain, the slur predications in [13] and [14] do not. They are linguistically aright, entirely felicitous.

Now, will these objections directly transfer to Identity Expressivism? Not straightaway. While the McCready and Potts accounts are both expressivist, they are in some important respects different from my own. The most fundamental is that both require that with a slur, speakers express some *negatively valenced attitude or affect*, not, in particular, contemptuous regard.¹⁷ With this in mind, we can construct a version of the Specificity Problem tailor-made for Identity Expressivism. It goes as follows:

Identity Expressivism's expressivist component is too confining because:

¹⁷ There is another important difference with the Potts analysis. Whereas Potts holds felicitous uses of an expressive modifier like *damn* and *fuck'n* requires the speaker be in a 'heightened emotional state right this minute', (Potts 2007: 171), that is a claim I explicitly reject – and for *both* expressives modifying nouns for social groups and for slurs. When combined with nouns as in *goddamn Jew* and *fuck'n homosexual*, the combined expression functions, on my view, exactly like slurs do. As we saw above, on Identity Expressivism, *Jake is a kike* is *synonymous* with *Jake is a goddamn Jew*. On this score, Stojnić and Lepore differential assessments of [8] and [8'] and [9] and [9'] seem to me exceptionally surprising and counterintuitive. We should register the same assessments about saying *You are a f****t!* and *You are a goddamn homosexual!* Following up either one with *I love faggots/homosexuals!* is certainly linguistically deviant. In what follows, I argue the full case for the slur.

- I. Competent users of slurs need not be occurrently experiencing feelings of contempt toward the group.
- II. Competent users of slurs need not have any standing feelings of contempt toward the group.
- III. Competent users of slurs may entirely lack all contempt or contemptuous attitudes toward the group.

In support of I-III, they can use the linguistic examples above. The rationale then becomes:

- IV. [9], [10], [11], [13], [14] manifest competent use. They are in no way linguistically incoherent, deviant, or infelicitous.

This is our construction of Stojnić and Lepore's *Specificity Problem for Identity Expressivism*. It is an important, radical, and powerful challenge.

The key to why this objection does not compromise Identity Expressivism is to understand the nature of contempt and the way it can be expressed in language. In previous work, I advanced a detailed analysis of contempt. A main impetus of the analysis was to address earlier, less radical versions of the Specificity Problem by Camp and Anderson and Lepore.¹⁸ There I drew heavily on two pools of research: one, empirical research by emotions theorists like Paul Ekman, Klaus Scherer, Ira Roseman, and June Tagneny; the other, philosophical analyses by moral psychologists like Michelle Mason, Macalaster Bell, and Steve Darwall.¹⁹ The following summarizes main points pertinent for combatting Stojnić and Lepore's Specificity Problem for Identity Expressivism. All of [i]-[vii] concern the *nature of contempt* itself.

[i] Contempt is a complex, though basic emotion, typically experienced as a *form of regard*. It is an *affective stance*.

[ii] Contempt is a hierarchizing affective stance. In harboring contempt for a person, one looks down on them, regarding them as low, as lesser.

[iii] Contemptuous regard must be sharply distinguished from contempt's *behavioral manifestations*. These can vary enormously,

¹⁸ Cf., Jeshion (2018). Camp's objection is that not all weapon-uses of slurs are associated with contempt. Although some are, she acknowledged, others are associated with "disgust, fear, dismissiveness" Camp (2013: 10). Further, one need not outwardly manifest contemptuous regard like hateful feeling. To underscore the point, she offered examples like this felicitous weapon-use of slurs: *I'm glad we have so many Sp*cs at our school; they always bring the best food*. Anderson and Lepore offered examples that register compliments toward the group: *Ch*nks are so much smarter than the rest of us*. I argued that Camp and Anderson and Lepore's challenges rest on false assumptions about the nature of contemptuous regard and its expression in language.

¹⁹ Cf., Ekman and Friesan (1986), Ekman and Heider (1988), Ekman (1994), Mason (2003), (2016), Scherer (2003), (2013), Tangney, J. Stuewig, J. & Mashek, D. (2007), Bell (2013), Roseman (2018).

ranging from manifestations of hate, demonstrations of disgust, mocking laughter, sneering, condescending pity, amused dismissiveness, simmering resentment, and simple indifference.

[iv] Because it is a standing form of regard, having contempt toward X at time t is compatible with the absence of behavioral manifestations of contempt toward X at time t.

[v] Contemptuous regard doesn't have a specific phenomenology. Unlike emotions like anger and fear, there is no specific way that it feels like to regard another with contempt. There is no well-defined 'feeling of contempt'.

[vi] Having contemptuous regard toward X is compatible with appreciating X's positive features.

[vii] Having contemptuous regard toward X is compatible with being self-blind to one's own contempt. People harbor contempt toward others without knowing they do.

From these features of the nature of contempt, we can straightforwardly derive three points about the *expression of contempt in language*.

[viii] One can felicitously use a term that expresses contemptuous regard without accompanying it with any behavioral manifestations of contempt.

[ix] One can felicitously use a term that expresses contemptuous regard without experiencing or communicating any alleged 'feelings of contempt'.

[x] Expressing contemptuous regard for X is compatible with sincerely communicating appreciation of X's positive features.

For deeper understanding of and support for each point, I refer the reader to the full analysis and especially the empirical and philosophical accounts that underwrite it. Even still, absent that, it is easy to recognize the truth of many of them. For instance, if you regard a certain president with contempt, you can do so *sans* behavioral manifestations, and without any specific phenomenology. Further, with respect to [vi] and [x], there is simply no problem with your expressing your contempt thus

[15] *Goddamn Trump* came into office in 2025 fantastically organized and focused,

while acknowledging (resentfully) positive properties.

These features of contempt straightaway undercut claims I and II of Stojnić and Lepore's Specificity Problem as a case against Identity Expressivism. Bigots can indeed competently use slurs without occurrently experiencing 'feelings' of contempt or manifesting any behaviors that signal contempt. Thus, Identity Expressivism is in fact compatible with I and II.

What about III, Stojnić and Lepore's claim that no negative attitude or affect or stance is necessary for competent use?²⁰ This is different. Identity Expressivism *requires* that competent speakers have contemptuous regard, a stance obviously incompatible with the total absence of standing negative attitudes, affects, feelings, and opinions about the target group. In support of III, Stojnić and Lepore offer a series of examples, consolidated here.

[9] I don't think C***ks are despicable.

[10] John is a n****r. N****rs are very respectable people.

[11] F****ts are good! We should respect them.

[13] Leo is a Kike. But, to be clear, I have no ill feelings toward kikes.

[14] You are a f****t. I love f****ts, they are wonderful people!

These uses of slurs are coupled with: denials of thinking the group despicable [9], avowals of respect [10] and directives for all to respect the group [11], denials of having any negative feelings [13], and expressions of reverence and love toward the group [14]. I take each in turn.

Unlike the rest, there is no problem with [9]: sincere, knowledgeable denial that one thinks the group despicable can be felicitously used with a slur, as discussed above in the Schlenker analysis. Furthermore, contemptuous regard only requires looking down on the group, seeing them as lesser; it does not require thinking they are *despicable*. Thus, [9] offers no support to III, as it does not show absence of all negative feeling. And by itself, [9] is fully compatible with Identity Expressivism.

By contrast, the avowals of respect in [10] and directives for all to respect the group in [11] strike one – strike me – as highly marked, at least if the respect entails regarding the group as *in no way* lesser. To appreciate this, remember that in testing for felicitousness, we must assume the speaker's use of the slur is a basic use, a weapon use. We must also assume that the speaker possesses full linguistic understanding of the slur. In the absence of these conditions, [10] and [11] might sound acceptable. Otherwise, they land as highly marked, smacking of a speaker ludicrously squirming toward some curious form of plausible deniability.

The same should be said for the wholesale denial of all negative feelings in [13] and avowals of love and reverence for the group in [14]. They immediately land as deviant, infelicitous. One wants to say: why are you derogating the group you (profess to) love? Why choose the slur, if you truly harbor no ill feelings? To even make sense of such a weapon use with full understanding, we lean toward ascribing some measure of irrationality, positing either a blindness to their own affects and feelings coupled with a smarmy form of insincerity.

²⁰ Cf., also Anderson and Lepore's claim that slur use is compatible with harboring "no negative opinions towards" the target group (2013).

A comparison with racially coded dogwhistles sheds light on this. Dogwhistles are communicative devices that enable speakers to send cloaked messages. Racially-coded dogwhistles are expressions like *inner city* that can be used to function to communicate racist attitudes. Unlike racial slurs which are universally recognized as blatant tool of derogation, words commonly used to dogwhistle have basic uses that are completely neutral, as in *The inner city kids need state-supported lunches*. Relatedly, they don't wear their racism on their sleeves. When used as dogwhistles, the racist message is somehow hidden – available only to those in the know, or at least thinly veiled. Take, for instance, the racist dogwhistle *inner city* in this selection from Donald Trump's *Time to Get Tough: Making American #1 Again*:

If we keep on this path, if we reelect Barack Obama, the American we leave our kids and grandkids won't look like the America we were blessed to grow up in. The American Dream will be in hock. The shining city on the hill will start to look like an inner-city wreck.

On influential analyses by Saul and Khoo, dogwhistles do not have their hidden meaning as part of their semantic content.²¹ *Inner city* is synonymous with *densely populated urban area*; its meaning is devoid of anything about race. Trump uses *inner city* to signal race because he knows something about his hearers beliefs: that the inner city “brings to mind poor, crime-ridden, African-American neighborhoods” (Khoo 2017: 34). Thus, while saying only, ‘innocently’, that a second Obama term will result in an America that looks like a densely populated urban area, he messages that it will result in “an America dominated by poor, lazy, and criminal African Americans” (Khoo 2017: 34).

Dogwhistlers often accompany their utterances with what Saul has aptly called *figleaves*, addendums to the utterance that function as a way to plead innocent. When someone calls Trump's remarks out as racist, he might say *I was talking about the troubles in cities. Anyway, I love Black people!* Read: no racism here. Literal figleaves, as on Michaelangelo's David, provide cover for artists whose artworks would otherwise be viewed as overly sexual. In similar fashion, racist figleaves provide some modicum of cover for dogwhistlers' racism (Saul 2024: chapter 1). Notably, the figleaf of avowing love for Black people isn't incompatible with use of *inner city*.

Now return to [14]. After saying *You are a f****t*, the bigot follows up with *I love f****ts!* One deeply derogates with the slur and in the next moment professes love. The latter statement wrecks of both linguistic incoherence and insincerity, as do each of the follow-up remarks in [10], [11], [13]. They make the most sense construed as speakers desperately, yet ludicrously, grabbing for figleaves. Of course, the follow ups cannot provide cover, but the practice of grabbing a figleaf has become so normalized, it is easy to imagine someone thinking it could. Unlike *Anyway, I love Black people!* as an addendum to the dogwhistle,

²¹ Cf., Saul (2018) and Khoo (2017).

an utterance neither linguistically incoherent nor infelicitous, professing love for the group while using the slur is blatantly both. All in all, it's deeply misguided to pass these utterances off as linguistically non-deviant. Therefore, none provide support for III. Our dialectic reveals that this Specificity Problem leaves Identity Expressivism uncompromised.

4 Conclusion: Looking forward

In *Inflammatory Language*, Stojnić and Lepore boldly widen the purview of phenomena for theories of slurs to explain. The widening, together with their commitment to a single mechanism view, impels them to advance the remarkable view that slurring words themselves are not the ultimate source of slurs' pejorative sting. This piece has been dominantly defensive, aimed at showing that even if we grant the centrality and importance of mere orthographic and acoustic resemblance cases, their negative thesis does not stand. The first half demonstrated that their Single Mechanism Necessity Assumption rules out multiple mechanism theories. The second half demonstrated that their Specificity Problem does not undermine Identity Expressivism. However, Stojnić and Lepore's widening does an important service in bringing to the forefront deep and critical questions: What is the proper scope of a theory of slurs – what phenomena must it explain? What phenomena, if any, should take priority? What starting assumptions shall we make? What are the proper methodological principles guiding our investigations? Currently, the literature on slurs is disunified. We need a reckoning on the full and diverse range of issues before us and a systematic meta-analysis of the most pressing linguistic, psychological, social, and moral problems comprehensive theories of slurs must address.

References

- Anderson, L. and Lepore, E. 2013. "Slurring Words." *Noûs* 47: 25–48.
- Bach, K. 2018 "Loaded Words: On the Semantics and Pragmatics of Slurs." In Sosa (ed.). *Bad Words*. Oxford: Oxford University Press, 60–76.
- Camp, E. 2013. "Slurring Perspectives" *Analytic Philosophy* 54 (3): 330–349.
- Camp, E. 2018. "A Dual Act Analysis of Slurs." In Sosa (ed.). *Bad Words*. Oxford: Oxford University Press, 29–59.
- Davis, C. and McCready, E. 2020. "The Instability of Slurs." *Grazer Philosophische Studien* 93 (3): 63–85.
- Diaz-Legaspe, Liu, and Stainton. 2019. "Slurs and Register: A Case Study in Meaning Pluralism." *Mind and Language* 35 (2): 156–182.
- Ekman, P. and Friesan, W. 1986. "A New Pan-Cultural Facial Expression of Emotion." *Motivation and Emotion* 10 (2): 159–168.
- Ekman, P. and Heider, K. 1988. "The Universality of Contempt Expression: A Replication" *Motivation and Emotion* 12: 303–308.

- Ekman, P. 1994. "Strong Evidence for Universals in Facial Expression: A Reply to Russell's Mistaken Critique." *Psychological Bulletin* 115: 268–287.
- Fasoli, F. A. Maass, and A. Carnaghi. 2015. "Labeling and Discrimination: Do Homophobic Epithets Undermine Fair Distribution of Resources?" *British Journal of Social Psychology* 54 (2): 383–393.
- Fischer, A. & Roseman, I. 2007. "Beat Them or Ban Them: The Characteristics and Social Functions of Anger and Contempt." *Journal of Personality and Social Psychology* 93: 103–115.
- 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.
- Hom, C. and May, R. 2015. "Pejoratives as Fiction." In Sosa (ed.). *Bad Words*. Oxford: Oxford University Press, 108–131.
- Jeshion, R. 2013a. "Slurs and Stereotypes." *Analytic Philosophy* 54 (3): 314–329.
- Jeshion, R. 2013b. "Expressivism and the Offensiveness of Slurs." *Philosophical Perspectives* 27: 231–259.
- Jeshion, R. 2016. "Slur Creation, Bigotry Formation: The Power of Expressivism" *Phenomenology and Mind* 1: 130–139.
- Jeshion, R. 2017. "Loaded Words, Expressive Words: Assessing Two Semantic Frameworks for Slurs" *Croatian Journal of Philosophy* 50: 111–130.
- Jeshion, R. 2018. "Slurs, Dehumanization, and the Expression of Contempt." In Sosa (ed.). *Bad Words*. Oxford: Oxford University Press, 77–107.
- Jeshion, R. 2021. "Varieties of Pejoratives." In Khoo and Sterkin (eds.). *The Routledge Handbook of Social and Political Philosophy of Language*. Routledge Press, 211–231.
- Jeshion, R. 2025. "Slurs, Articulations, and Inflammatory Language" forthcoming. In Sosa and Lepore (eds.). *Oxford Studies in Philosophy of Language*, volume 4.
- Jorgensen Bolinger, R. 2017. "The Pragmatics of Slurs." *Noûs* 51 (3): 439–462.
- Khoo, J. 2017. "Code Words in Political Discourse." *Philosophical Topics* 45 (2): 33–64.
- Lepore, E. and Stone, M. 2014. *Imagination and Convention*. Oxford: Oxford University Press.
- Marques, T. and Garcia-Carpintero, M. 2020. "Really Expressive Presuppositions and How to Block Them." *Grazer Philosophische Studien* 97 (1): 138–158.
- Mason, M. 2003. "Contempt as a Moral Attitude." *Ethics* 113: 234–272.
- Mason, M. 2016. "Contempt at the Limits of Reactivity." In Mason, ed., *The Moral Psychology of Contempt*. London: Rowman & Littlefield Publishers, 173–192.
- Nunberg, G. 2018. "The Social Life of Slurs." In Fogel, Harris, & Moss (eds.). *New Work on Speech Acts*. Oxford University Press, 237–295.
- Potts, C. 2005. *The Logic of Conventional Implicatures*. Oxford: Oxford University Press.

- Potts, C. 2007. "The Expressive Dimension." *Theoretical Linguistics* 33 (2): 165–197.
- Rappaport, J. 2019. "Communicating With Slurs." *The Philosophical Quarterly* 69: 795–816.
- Richard, M. 2008. *When Truth Gives Out*. Oxford: Oxford University Press.
- Rinner, S. and Hieke, A. 2021. "Slurs Under Quotation." *Philosophical Studies* 1483–1494.
- Roseman, I. 2018. "Rejecting the Unworthy: The Causes, Components, and Consequences of Contempt." In Mason (ed.). *The Moral Psychology of Contempt*, 107–130.
- Saul, J. 2018. "Dogwhistles, Political Manipulation, and Philosophy of Language." In Fogal, D., Harris, D., Moss, M. (eds.). *New Work on Speech Acts*. Oxford University Press, 360–383.
- Saul, J. 2024. *Dogwhistles and Figleaves: How Manipulative Language Spreads Racism and Falsehood*. Oxford: Oxford University Press.
- Scherer, K. 2003. "Vocal Communication of Emotion: A Review of Research Paradigms." *Speech Communication* 40: 227–256.
- Scherer, K. 2013. "Emotion in Action, Interaction, Music, and Speech." In Arbib (ed.). *Language, Music, and the Brain: A Mysterious Relationship*. MIT Press, 107–139.
- Schlenker, P. 2007. "Expressive Presuppositions." *Theoretical Linguistics* 33 (2): 237–245.
- Stillman, R. 2021. "Slurs as Ballistic Speech." *Synthese* 199: 6827–6843.
- Stojnić, U. and Lepore, U. 2025. *Inflammatory Language*, final manuscript version.
- Tangney, J. Stuewig, J. and Mashek, D. 2007. "Moral Emotions and Moral Behavior." *Annual Review of Psychology* 58 (10): 345–372.