

Do laudatives really mirror pejoratives?

Pepa Mellema

Date: 14/08/2020

Student number: s1019393

Supervisor: prof. dr. L.B.W. Geurts

Word count (excluding references): 10.172

Thesis for obtaining a “Master of arts” degree in philosophy
Radboud University Nijmegen

I hereby declare and assure that I, Pepa Mellema, have drafted this thesis independently, that no other sources and/or means other than those mentioned have been used and that the passages of which the text content or meaning originates in other words—including electronic media—have been identified and the sources clearly stated. Place: Nijmegen, date: 14/08/2020

Abstract

Pejoratives are words like *jerk* that convey negative attitudes towards their targets. Laudatives are words like *sweetheart* that convey positive attitudes towards their targets. Some theorists have hypothesised that laudatives are merely positively-valenced analogues to pejoratives, but have not provided empirical evidence to support this hypothesis. The present paper therefore serves two goals: (i) to document the (dis)similarities between laudatives and pejoratives and (ii) to analyse the theoretical consequences of these findings for the philosophical debate about derogatory content. I argue that laudatives largely mirror the behaviours of pejoratives and conclude that non-content-based approaches to pejoratives are better at explaining the minor differences between the two classes than are content-based approaches.

1 Introduction

Pejoratives are words that convey derogatory attitudes towards their targets. We expect competent speakers of English to know that to call someone a *bastard* or *faggot* is to derogate them. We also expect competent speakers to know that using pejoratives is often objectionable. To call someone a *dickhead*, or worse, a *faggot*, violates social norms that prohibit speakers from using these words. These observations have prompted a chicken-and-egg dilemma in the philosophy of language. What was there first: the prohibitions on using pejoratives or the pejoratives' capacity to offend? Do we prohibit words like *dickhead* and *faggot* because they mean something bad, or do they mean something bad because we prohibit them?

Taking the first horn of the dilemma, content-theorists argue that we ban pejoratives because their linguistic contents are objectionable. Although there is no consensus on what comprises derogatory content, proponents agree that pejoratives are offensive because of their contents. Some advocates of this approach are [Hom & May \(2013\)](#), [Camp \(2013\)](#) and [Jeshion \(2013\)](#). Taking the second horn of the dilemma, proponents of non-content-based theories argue that there is nothing objectionable about the contents of objectionable words. On this view, our judgements that *bastard* and *faggot* are offensive are akin to our judgements that *homie* and *dude* are informal, or that *betwixt* and *forthwith* are archaic. These judgements do not reflect the hearer's understanding of a word's contents. Rather, these judgements reflect the hearer's knowledge of sociocultural facts that go beyond the word's

contents. To know a word is offensive is to know its use is objectionable or forbidden. Proponents of this approach are [Anderson & Lepore \(2013\)](#), [Lepore & Stone \(2018\)](#) and [Pullum \(2018\)](#).

In the present article, I argue that we can break the gridlock between these opposing views by investigating a different class of evaluative words altogether—the so-called laudatives. Paradigmatic examples of laudatives are *saint*, *stud*, and *genius*. Like pejoratives, these words convey evaluative attitudes towards their targets. But unlike pejoratives, laudatives convey positive laudatory attitudes. To call someone a *saint* or a *stud* is to laud them.

In the philosophical literature on evaluative words, laudatives have been largely ignored except for the occasional comparison to pejoratives. For instance, [Nunberg \(2018\)](#) and [Mišćević \(2017\)](#) have hypothesised that laudatives are merely positively-valenced analogues to pejoratives. However, neither Nunberg nor Mišćević have systematically compared the linguistic behaviours of laudatives against those of pejoratives in order to check their hypothesis. This, I think, is a missed opportunity since any (dis)similarities between the two classes could inform and constrain theorising about derogatory language. The present paper therefore serves two aims: (i) to document the extent to which laudatives mirror pejoratives and (ii) to analyse the theoretical consequences of these findings for the philosophical debate about derogatory content.

My plan is the following. I first identify several key characteristics of pejoratives (section 2) before surveying the recent debate on whether these characteristics can be explained in terms of linguistic contents (section 3). Using insights from the first two sections, I then compare the behaviours of laudatives against several well-documented behaviours of pejoratives. To conclude, I explore the consequences of the (dis)similarities between laudatives and pejoratives for the philosophical debate on derogatory linguistic content (section 5).

I will conclude that [Nunberg \(2018\)](#) and [Mišćević \(2017\)](#) were right to speculate that laudatives closely resemble pejoratives. However, I will also argue that there are two features of pejoratives that, although not completely lacking in laudatives, are exceedingly rare—that is, laudatives rarely offend and rarely project. I will conclude that content-theorists struggle to account for these minor dissimilarities, whereas their detractors fare better.

2 What pejoratives do

In order to compare laudatives to pejoratives, we will need to know what kind of words pejoratives are. The present section therefore surveys the linguistic features that unite pejoratives into a single group and the additional linguistic features that divide them into separate subgroups. What follows below is by no means an exhaustive taxonomy; it's a rough sketch that serves my description of laudatives later on.

2.1 Derogating

Let's begin with the feature that unites pejoratives into a single group—at least, nominally—namely, the capacity to convey derogatory attitudes towards a target. This capacity is commonly known as a word's pejorative or derogatory force. Here I use the terms *pejorative* and *derogatory* interchangeably. Before explaining what it means for a word to convey derogatory attitudes, let me first explain what it does not mean. More specifically, derogating should not be confused with offending hearers or with expressing emotions.

For a start, a common assumption in the literature on pejoratives is that derogating differs from taking offence¹. On this assumption, a word *derogates* insofar it conveys a derogatory attitude. By contrast, a word *offends* insofar a hearer experiences psychological distress at interpreting a word. In other words, derogating describes what words do; it is a feature of derogatory words. Taking offence is what hearers do; it describes their negative subjective experience of hearing a word.

The main reason to adopt this distinction is that there is a double dissociation between a word conveying a derogatory attitude and a hearer taking offence at the attitudes conveyed by the word. On the one hand, it is possible to use a derogatory word without causing offence. For instance, not everyone who gets called a *dickhead* or a *faggot* will feel offended. On the other hand, it is possible to use a word that does not typically express derogatory attitudes, but still cause offence to hearers. For example, the words *girl* and *grandpa* do not express negative attitudes by default, yet sometimes hearers take offence at being called a *girl* or *grandpa*. So, derogating is not offending.

Derogating should also not be mistaken for the expression of emotions. However, talk of derogation is often indistinguishable from talk of expressing strong negative emotions such as hatred and contempt. To call, say, Bob a *bastard* is to express moderate dislike of him. And to call him a *faggot* is to express sizeable contempt towards

¹E.g., see [Hom \(2012\)](#) or [Anderson & Lepore \(2013\)](#).

him. Representative here is [Hom \(2012\)](#) who writes: “These words allow speakers to convey *emotional states* beyond the truth-conditional contents that they are normally taken to encode.” (emphasis added, p.383).

I find this way of talking about derogation somewhat misleading. It conflates a word’s ability to emote with its ability to evaluate. But it is a mistake to lump together the expression of affect with the expression of evaluation: emotions and evaluations are different things. In this vein, [Rappaport \(2019\)](#) distinguishes between a word’s affective and evaluative capacities. He characterises a word as *evaluative* insofar it is “used to communicate the speaker’s negative attitudes” (p.795). He characterises a word as *affective* insofar it is “capable of expressing powerful emotions and causing a strong emotional response in hearers.” (idem.). I recommend we respect Rappaport’s evaluative–affective distinction when describing pejoratives and laudatives.

However, Rappaport doesn’t specify what an evaluation consists in, and how (if at all) it differs from affect. This is unfortunate since, although emotions and evaluations appear conceptually distinct, they do often coincide. It is perfectly natural to assign valences to emotions: angry is negative, happy is positive. In order to capture this observation—and, in order to aid my description of pejoratives and laudatives—let’s take a brief detour to consider what evaluative attitudes are, and how they relate to affective attitudes.

2.2 Evaluating

I suggest we take an *evaluative attitude* to be an attitude that is both (a) valenced and (b) targeted. An attitude is *valenced* insofar it has a valence assigned to it—either a positive or negative valence, or both simultaneously (bivalent). Valences may vary in strength: we might evaluate a target as being anywhere between extremely positive and extremely negative. An attitude is *targeted* insofar its valence is directed towards some contextually salient target. Targets may vary in kind and number: they may be hypothetical, real or imagined. A target may be a single object or individual, or a grouping of objects or individuals. An attitude may also target events or states of affairs.

This conception does not presuppose anything about how an evaluative attitude is cognitively realised. A single evaluative attitude might be realised by a single belief or by several related beliefs, desires and/or feelings simultaneously. For all I am concerned, an agent could have an evaluative attitude without having any mental states; it might take behaviours, dispositions, moral judgements, or whatever else takes

the reader’s fancy. All that is required for an attitude to qualify as evaluative is it being valenced towards some salient target. This may entail disliking Marmite or believing fascists to be evil.

The difference between *evaluative* and *affective* on my view is thus this: whereas the label *evaluative* is neutral about the underlying cognitive machinery required for having an attitude, the label *affective* identifies a specific subset of mental states responsible for realising the attitude, namely: emotions. This view does justice to the intuition that emotions and evaluations are related yet conceptually distinct. After all, most emotions neatly fit the mould of a valenced/targeted attitude. If you’re angry because you’ve lost your keys then you have a negative evaluative attitude towards having lost your keys.

To return to my discussion of derogatory force: when I say all pejoratives communicate derogatory attitudes, I am saying that pejoratives communicate negative evaluative attitudes in the sense defined above. On this view, the class of pejoratives includes all words that can be reasonably said to communicate negative evaluative attitudes. The resulting class of words roughly encompasses the three subgroups of pejoratives identified by Hom (2010): (1) syntactically flexible swearwords such as *fucking*, *rotten* and *damn*; (2) relatively mild terms of abuse such as *freeloader*, *pencil pusher* and *dickhead*; and (3) deeply offensive slurs such as the n-word, *faggot* and *kike*. This is clearly a heterogeneous group. Let’s therefore now consider three additional features that account for the variation amongst pejoratives. More specifically, some—but not all—pejoratives can describe, slur and project.

2.3 Describing

For a start, some pejoratives describe their targets. A word describes its target when it predicates its referent as belonging to a category of objects. Some words are purely descriptive. For example, *woman*, *combine harvester* and *tetrahedron* pick out descriptive features without expressing evaluative attitudes. To describe an object as a *tetrahedron* categorises it as a four-faced polyhedron, but does not indicate whether the speaker is partial towards four-faced polyhedra.

In contrast, most pejoratives are hybrids in the sense that they describe as well as evaluate their targets. To butcher Potts’ (2007, p.168) infamous example: calling Kresge a *bastard* negatively evaluates Kresge and describes him as socially maladaptive. The same goes for slurs: calling Bob a *faggot* characterises Bob as gay while derogating

him for this characterisation. Barring a few exceptions², theorists typically flesh out this capacity in terms of a word’s capacity to convey propositional content. On this approach, hybrid pejoratives describe their targets through conveying truth-evaluative content about them.

Not all pejoratives convey descriptive content. Interjections like *Bloody hell!* and *Fuck!* do not contribute to the truth-evaluable content of the sentences they occur in. For instance, the sentences “Snuggles died.” and “Fuck, Snuggles died.” are logically equivalent. Pejorative interjections like *Fuck!* instead serve to transmit the speaker’s evaluative attitudes towards some contextually salient target.

That is not to say that hybrid pejoratives always contribute truth-evaluable content. A well-documented feature of hybrid pejoratives is that speakers may felicitously use them to either derogate targets without accurately describing them; or, to describe targets without necessarily derogating them³. For example, speakers can successfully derogate heterosexuals using homophobic slurs. Conversely, speakers may felicitously use a hybrid pejorative to describe someone without necessarily invoking the pejorative’s derogatory effects. This can happen when the pejorative’s intended targets use the words amongst themselves as a term of endearment.

2.4 Slurring

A common refrain in the literature is that we should also divide the class of pejoratives into slurring and non-slurring ones. Representative examples of slurring pejoratives (or, *slurs* for short) are the n-word, *faggot* and *kike*. Representative examples of non-slurring pejoratives are *bastard*, *dickhead* and *asshole*. These subgroups behave similarly enough; being equally descriptive, evaluative and affective. At the same time, they appear qualitatively different. Let’s consider two suggestions for telling apart slurs from garden-variety pejoratives.

According to the social group criterion, slurs differ from other terms of abuse in virtue of derogating a different type of target. For example, [Blakemore \(2015\)](#) opts for this criterion. Whereas slurring pejoratives derogate individuals on the basis of belonging to a social group, non-slurring pejoratives derogate individuals on the basis of possessing some personal trait. If we accept this criterion, then *faggot* and *kike* are slurs since they derogate individuals on the basis of social group memberships. And *bastard* and *dickhead* are not slurs since they derogate individuals on

²One such exception is [Richard \(2008\)](#), who maintains that slurs cause truth-gaps. The truth of sentences containing slurs can therefore supposedly not be evaluated.

³See [Croom \(2013\)](#) and [Jeshion \(2013\)](#) for further discussion.

the basis of perceived individual shortcomings. The underlying reasoning here is that we don't typically think of socially maladaptive 'dickheads' as constituting a distinct social group like Jewish or gay people do.

Social group membership is clearly important to differentiating slurring pejoratives from non-slurring ones. However, as various authors have noted, the social group criterion is too inclusive since not all words that derogate individuals on the basis of social group memberships are slurs⁴. Consider *snowflake* for easily-offended millennials, *latte liberal* for hypocritical progressives, and *honky*, *cracker* or *white trash* for white people. The social group criterion would peg these as slurs. But not everyone would agree with this assessment. Some hearers might claim that there is an intuitive difference between being called the n-word and being called a *honky* or *cracker*. According to these hearers, the relevant difference is that slurs against African Americans are tied to a long history of racism, violence and oppression in ways that 'slurs' against white Americans are not. Yet, other hearers might be offended by the suggestion that *honky* is not a proper slur, countering that white Americans are subject to racism, too.

Even though I think there is little merit to the latter position, important here is why hearers disagree about the slurhood of *honky* and *cracker*: their disagreement is not about the meaning of the words themselves—rather, their disagreement stems from differing beliefs about which racial groups are victims of racism and which racial groups are not. In other words, hearers can disagree about which pejoratives are slurs and these disagreements derive from differing world views. Ideally, a definition of slurhood should capture these observations.

Nunberg (2018) articulates such a definition—and one that I will adopt here. His core insight is that slurs are not a linguistic, but a social kind. To categorise a word as a pejorative, suggests Nunberg, is to pass a linguistic judgement: it is to recognise that a word conveys an evaluative attitude. By contrast, to label a pejorative a slur is to pass a normative judgement: calling a pejorative a slur condemns the derogation and oppression associated with its use.

Poignantly, Nunberg observes:

[...] a slur is a kind of verbalized thoughtcrime: it perpetuates social inequities, infects even innocent minds, and undermines the conduct of public discourse. (2018, p.239)

On this view, a hearer's judgement about the slurhood of a word, reflects the hearer's beliefs about which (groups of) individuals deserve

⁴See Anderson & Lepore (2013), Popa-Wyatt (2016), and Nunberg (2018).

derogation, and which don't. For instance, when hearers label the n-word a slur, they condemn it for unjustly derogating individuals for being African Americans. Similarly, when hearers label *honky* and *cracker* slurs, they condemn the words for unfairly derogating individuals for being white Americans.

A consequence of this view is that there is no principled linguistic distinction between slurring and non-slurring pejoratives: whether a word is a slur is in the eye of the beholder. The distinction between slurring and non-slurring pejoratives is therefore as blurry as any individual hearer's understanding of what groups merit social justice, and what groups don't. This explains why slurs are sometimes defined as pejoratives that target historically marginalised groups such as the LGBTQ+ community or racial, ethnic and religious minorities. This is because most people—but again, not all—agree that these groups are unjustly marginalised and should therefore not be subject to derogation.

2.5 Projecting

Another source of variation amongst pejoratives is that some, but not all, pejoratives exhibit unusual projective behaviour. Theorists commonly characterise projection as follows: a pejorative *projects* if it remains offensive and/or derogatory despite being in the scope of functional embeddings such as truth-conditional and intensional operators⁵. Particularly prone to projecting are slurring pejoratives like *faggot* and *kike*, and expletive pejoratives like *dickhead* and *asshole*:

- (1)
 - a. **negation**: Bob's not a faggot
 - b. **conditional**: If they cheat, they're dickheads
 - c. **counterfactual**: Bob isn't Jewish, but if he had been, he'd be a kike
 - d. **tense**: Ann used to be a cunt, but now she's nice

This insensitivity supposedly extends to reported pejoratives:

- (2)
 - a. **speech report**: Bob said "Ann is a cunt"
 - b. **mental state report**: Ann soon realises Bob is an asshole

By contrast, non-slurring pejoratives such as *latte liberal* and *pencil pusher* don't project. Instead, their derogatory force interacts with

⁵E.g., Bolinger (2017); Hom (2008, 2012, 2020); Potts (2007).

operators in the way we expect propositional content to interact with operators⁶:

- (3) a. **negation**: Bob’s not a jerk
- b. **conditional**: If Bob cheats, then he’s a deadbeat
- c. **counterfactual**: Ann didn’t vote for Hillary, but if she had, she’d be a latte liberal
- d. **tense**: Ann used to be a commie, but now she votes conservative

The literature on pejoratives does not provide a satisfactory description of these data. Some theorists just note that some pejoratives are insensitive to modification by functional embeddings, but stop short of specifying what aspect of the pejoratives’ meaning avoids modification. Other theorists have tried to identify what projects out, but in doing so often presuppose specific theories of pejorative force, or worse, in doing so ignore distinctions such as those between offending and derogating; or between evaluating and emoting.

Take Hom’s discussion of projection. According to Hom, pejoratives project when their “pejorative *content* appears to scope out; that is, even in nonassertoric contexts, slurs are deeply *offensive*.” (2020, p.289, emphasis added). “To be more precise”, Hom adds: “slurs avoid functional modification relative to various truth-functional and intensional operators that have them in their scope (idem.)”. Here Hom muddles the distinction between offending and derogating. He also makes his description of projection theory-laden by presupposing that derogatory attitudes are encoded into the propositional content of pejoratives.

Anderson & Lepore’s (2013) discussion of projection is equally messy. Like Hom (2020), they use the labels *derogatory* and *offensive* interchangeably. On top of that, they provide several distinct characterisations of projection. First, they identify projection with what happens when embedding a slur “inside a sentence does not immunize its users from transgression, even though sentential embedding can render semantic (as well as pragmatic) properties inert (p.353).” Later they characterise projection as slurs “refuse[ing] to submit to a unilateral detachment of “the affect, hatred and negative connotations

⁶Note that hybrid pejoratives are usually ambiguous between evaluative and descriptive readings. This ambiguity persists when hybrid pejoratives are modified by truth-functional operators. For example, “Bob’s *not* a deadbeat” yields an evaluative reading (i.e., “I do *not* negatively evaluate Bob”) and a descriptive reading (i.e., “Bob is *not* an unproductive member of society”). Here I am only considering evaluative readings.

tied to most slurs⁷” (2013, pp.353–354). Here they identify projection with the scoping out of both evaluative and affective attitudes, while also equating it with the offence generated by violating prohibitions against the use of pejoratives.

So, although it is easy to recognise a projecting pejorative, it is hard to articulate what part of its meaning escapes functional embedding. Let’s briefly consider some possible ‘escapees’. To begin, I think we can immediately eliminate the propositional content of pejoratives as a plausible candidate. This is because the propositional content of projecting pejoratives behaves as the propositional content of regular descriptive terms. This is best explained with an example:

- (4) a. Ann is a cunt
b. Ann is a mechanic

These sentences are identical except that sentence (4a) contains the hybrid pejorative *cunt*, whereas (4b) contains the regular descriptive term *mechanic*. Yet conditionalising or negating these sentences yields sentences with the same truth-conditions. For instance, we can make the truth of both sentences dependent on the truth of an antecedent:

- (5) a. If Ann killed Snuggles, then she is a cunt
b. If Ann fixed the car, then she is a mechanic

Here the propositional content of the pejorative *cunt* responds to the if-then construction in the same way as the propositional content of the regular descriptive term *mechanic*. So, when pejoratives project, it is probably not their propositional content that escapes embedding.

More plausible is that what projects is the pejorative’s derogatory, or negative evaluative, attitudes. This would explain why slurs appear to derogate everyone who could be targeted by their use, rather than the specific individuals who are being targeted by specific uses. For instance, both “Bob’s a faggot” and “Bob’s *not* a faggot” derogate everyone who is gay, including Bob. Even the sentence “I don’t object to homosexuality, some of my best friends are faggots” manages to derogate all gay individuals despite the speaker’s explicit disavowal of homophobia. This suggests that the derogatory attitudes of some slurs successfully avoid functional embedding.

Still, this does not account for all the projective behaviours of pejoratives: some non-slurring pejoratives project even though their derogatory attitudes do not scope out. Consider expletive non-slurring

⁷The quotation within the quotation is due to Richard (2008, p.62).

pejoratives such as *dickhead* and *asshole*. These clearly project in the sense that they remain offensive in all occurrences. However, if Ann calls Bob a *dickhead*, then she does not imply having derogatory attitudes towards *all* socially maladaptive individuals. She only signals her derogatory attitude towards Bob.

But, then, what scopes out when expletives project, if not their derogatory attitudes? My best guess is that, in these cases, the hearer simply takes offence at hearing an inappropriate word. I think [Anderson & Lepore \(2013\)](#) are onto something when they describe projection as the transgression of using a forbidden word. If hearers feel offended at hearing words that are forbidden, then these words will offend hearers in all occurrences. So, it is not the expletive's derogatory attitude that resists embedding—it's the psychological discomfort hearers experience when they encounter inappropriate words.

In any case, not all pejoratives project equally. Some pejoratives don't project, some do. Moreover, different aspects of pejoratives can project: slurs project derogatory attitudes and expletives project offensiveness. Some particularly egregious slurs probably project both their derogatory attitudes and their offensiveness.

2.6 Pejoratives summarised

In sum, pejoratives are not a uniform class of words. Although they all express negative evaluative attitudes towards their targets, they vary on a number of dimensions.

First, they vary with regards to how descriptive they are: some pejoratives descriptively identify their targets whereas others don't. Second, pejoratives vary with regards to how affective they are: speakers may, or may not, use pejoratives to express and/or evoke emotions. Third, pejoratives differ with regards to their social acceptability: some pejoratives derogate individuals who hearers find deserving of contempt, whereas others do not. The less acceptable hearers judge the attitudes communicated by a pejorative, the more likely they are to label it a slur. Fourth and last, pejoratives vary with regards to whether they project. Whereas non-slurring and non-expletive pejoratives don't project, expletive and slurring pejoratives are offensive in all instances of use.

3 Pejoratives and their (dis)contents

3.1 The debate

This peculiar collection of traits has fuelled decades of theorising about pejoratives. The core point of contention amongst philosophers of language concerns the nature of the mechanisms by which pejoratives derogate their targets. I plan to show that laudatives can inform and constrain such theorising. In order to argue for this position later, in this section I first examine several prominent theories meant to explain the derogatory capacity of pejoratives.

As [Lepore & Stone \(2018\)](#) and [Pullum \(2018\)](#) observe, theorists are divided into two camps: the content-theorists who argue that pejoratives derogate in virtue of possessing conventionally derogatory linguistic contents and the non-content-theorists who reject this view.

There are many different way to encode derogatory attitudes into the contents of pejoratives. Before turning to the non-content-based theories, let's consider three of them.

3.2 Content-based theories

One option is to encode the derogatory attitudes straight into the at-issue content. This position, known as semantic externalism, is favoured by [Hom \(2008, 2012, 2020\)](#) and [Hom & May \(2013\)](#). Its founding assumption is that a pejorative's derogatory force is fully reducible to its literal at-issue content. For instance, the at-issue content of the slur *faggot* conveys the complex predicate, “a gay person who is contemptible for being gay”, along with the unsavoury propositions that homophobes take to be true about gay people, such as “gay people are sinful, dangerous and unnatural”.

However, this view runs into considerable problems. Its most serious problem being that it comes packaged with the deeply counter-intuitive claim that pejoratives have empty extensions. No individual can actually instantiate the derogatory predicates supposedly encoded by pejoratives. For example, most of us believe that gay people are not contemptible, sinful, unnatural or dangerous for being gay. If we also accept semantic externalism, then no one can be a *faggot* or *dyke*. However, as [Sennet & Copp \(2014\)](#) point out: this is approach is noble but questionable since it yields bizarre consequences for the truth-conditions of sentences with pejoratives.

For a start, it makes numerous sentences containing pejoratives false, including sentence we might judge to be objectionable but true, such as “Bob's not straight, he's a faggot.” or “Ann killed a cat, what a

murderous cunt!” Even stranger: the null extensionality of pejoratives makes sentences such as “All faggots are straight” and “All kikes are Mormons” true by default, even though they are clearly false. Semantic externalism therefore isn’t a particularly plausible rendering of the semantics of pejoratives.

A second option is to posit an additional type of conventional meaning distinct from propositional content and use it to explain why pejoratives derogate. This is the strategy favoured by proponents of expressivism. Following [Kaplan’s \(1999\)](#) suggestion, theories of this pedigree postulate the existence of two types of linguistic content: (i) truth-evaluative propositional content and (ii) non-truth-evaluative expressive content that conventionally indicates the speaker’s “heightened emotional state” ([Potts, 2007](#), p.173). Words like *ouch* and *oops* are ‘pure’ expressives in that they only communicate the speakers’ affective states. According to expressivists, derogatory content is a species of expressive content. Pejoratives derogate their targets in virtue of expressing strong negative emotions towards them such as hatred and contempt. To call, say, Kresge a *bastard* is to express contempt towards him.

The shortcomings of expressivism are too many to enumerate here⁸. For now, it suffices to say that it is intuitively implausible that pejorative force is fully reducible to a word’s capacity to convey and/or evoke emotions. Case in point: empirical research bears out that emoting is not even the primary function of ‘bad’ words—not even of swearwords. Although the expression of “heightened emotional states” is undeniably a function of swearing, it has various non-affective yet equally important functions in interpersonal communication. [Stapleton \(2010\)](#) surveys the relevant literature and catalogues these interpersonal uses of swearing as follow: enhancing a joke or story, emphasising a point, establishing solidarity or intimacy, covering up vulnerability, constructing and displaying one’s identity, and threatening one’s audience. So, the expressivist’s suggestion that emotional expression is the core business of pejoratives is an unhelpful simplification.

A third strand of content-based theories opts for a lean semantics instead. According to these theories, hearers infer the speaker’s derogatory attitudes from the conventional meaning of pejoratives. There are several ways to go about this. For sake of brevity, let’s consider only one. One suggestion is that pejoratives communicate derogatory attitudes in virtue of triggering derogatory presuppositions.

Presuppositions are the kind of assumptions we need to interpret

⁸See [Blakemore \(2015\)](#), [Croom \(2014\)](#), [Geurts \(2007\)](#), [Jeshion \(2013\)](#), [Hom \(2012\)](#) and [Nunberg \(2018\)](#).

the at-issue propositional content of an utterance. For instance, suppose we intend to assess the truth of “Snuggles has perished”. Its truth can only be evaluated insofar we assume Snuggles to have been previously alive. Some suggest that presuppositions are also responsible for the derogatory capacity of pejoratives (e.g., [Cepollaro, 2015](#); [Schlenker, 2007](#)). For instance, to interpret “Bob’s a *faggot*” we need to presuppose homosexuality to be contemptible, unnatural and harmful.

However, as [Nunberg \(2018, p.64\)](#) objects, this analysis wrongly predicts that various sentences containing pejoratives should be tautological, when they are not. On a presuppositional treatment of the slur *faggot*, for instance, hearers can only interpret utterances containing *faggot* if they accept the presupposition that gays are harmful, unnatural and sinful. The sentence “Faggots are sinful” should then be as tautological as “Murderers are murderous” or “Ann has killed again and she has killed before”, but it is not. The sentence “Faggots are sinful” introduces new information. Or, to generalise to all pejoratives: hearers who know the conventional meaning of a pejorative can still learn new information when speakers use the pejorative to attribute stereotypical features to its target. It is therefore unlikely that presuppositions are fully responsible for delivering the derogatory attitudes of pejoratives.

The theories surveyed here represent only a small fraction of the available content-based theories. Notwithstanding the internal disagreements, content-theorists agree that understanding pejorative force ultimately is understanding how conventional linguistic meaning encodes derogatory attitudes.

3.3 Non-content-based theories

In recent years, a competing non-content-based approach has been gaining popularity. Advocates of this approach contend that pejorative terms pack the same linguistic contents as non-pejorative terms. The derogatory effects of pejoratives, they argue, should be explained in terms of metadata—not linguistic contents. A word’s metadata consist of sociocultural facts about the word that go beyond its linguistic contents⁹. On this view, our judgements that words like *bastard* and *faggot* are derogatory are akin to our judgements that *homie* and *dude* are informal, or that *betwixt* and *forthwith* are archaic.

Consider the differences in connotation between *dude* and *man*, or between *homosexual* and *faggot*. According to non-content-theorists these pairs carry the same linguistic contents: *man–dude* predicate a person as

⁹See [Pullum \(2018, pp.169–170\)](#) for further discussion of metadata.

male; *homosexual*–*faggot* predicate a man as gay. The differences between the paired words are differences in metadata—differences in sociocultural information about the word’s use pertaining to who typically uses the word, in what social contexts they use it and how appropriate such uses are. The word *dude* differs from the word *man* only in that it is an informal, chiefly American slang term that may be used as a term of address. Similarly, the slur *faggot* differs from the word *homosexual* in that it is a derogatory term, chiefly used by homophobes; associated with a long history of silencing, marginalising and dehumanising gay people.

Different iterations of non-content-based accounts highlight different types of metadata to explain pejorative force. Nunberg (2018) singles out metadata about which social groups commonly use which words. Anderson & Lepore (2013) spotlight metadata about which words are prohibited. And Pullum (2018) suggests that all types of metadata can contribute to derogatory force. Here I will only discuss, Anderson & Lepore’s (2013) view, better known as prohibitionism.

According to prohibitionists, a word is derogatory insofar its use is forbidden. To know *dickhead* or *faggot* are derogatory is to know there exist social norms that forbid speakers from using these words. Derogatory terms, on this view, only differ from co-referential descriptive terms in being taboo.

Camp (2013, p.343) identifies two key problems with prohibitionism. Firstly, a word’s taboo status alone cannot account for the differences between slurs and swearwords. Slurs and swearwords are both forbidden, yet the offence of using the n-word greatly trumps the offence of dropping a f-bomb. Prohibitions alone cannot explain these differences.

Secondly, prohibitionists struggle to explain why we ban some words rather than others if not for differences in their contents. If *faggot* and *homosexual* really have the same linguistic contents, then why would we ban one but not the other? Prohibitionism, objects Camp, reverses horse and cart: we ban slurs because they conventionally convey derogatory attitudes—they do not conventionally convey derogatory attitudes because we ban them. In other words, prohibitionists, fail to specify a mechanism by which pejoratives deliver the evaluative attitudes that hearers deem offensive.

To meet these concerns, Lepore & Stone (2018) have supplemented prohibitionism with a tonal theory, which subsumes the derogatory effects of pejoratives under tonal effects. Tone, they explain, is “a catchall description of interpretive effects that go beyond meaning in language: effects that are heterogeneous in origin, open-ended, and often non-propositional” (p.133). Tone *goes beyond meaning* in that tone concerns interpretative effects that are not part of the conventional

and/or propositional content of an utterance. Tone is *open-ended* in that the interpretation of an utterance is not an objective immutable feature of an utterance’s linguistic meaning. Instead, it depends on variable features of hearers such as their (dis)likes, beliefs and/or prior experiences that shade the hearers’ interpretations of a word’s linguistic meaning. The origins of tone are further *heterogeneous* in that they have multiple sources; no single cognitive, social or linguistic mechanism is solely responsible for tone.

On this amended view, our reasons for banning words may differ. We ban slurs because they convey a contemptuous tone towards social groups while we judge this tone to be unjust; or worse, because we associate it with systemic oppression. We ban swearwords for less lofty reasons: not for perpetrating social injustice, but for mentioning “unmentionables”—taboo topics such as bodily excretions, sexual organs or blasphemous acts. Still, prohibitionists maintain: both swearwords and slurs are derogatory on account of violating prohibitions against their use—not on account of their linguistic contents.

3.4 Why laudatives matter to this debate

The non-content-based approach radically departs from the content-based approach we saw earlier. Ultimately, determining which approach is better depends on its capacity to adequately explain the peculiar linguistic behaviours of pejoratives. I believe that a better understanding of laudatives may aid theory choice, too. This is because, as I will argue in the next section, laudatives are similar enough to pejoratives to suspect that their shared behaviours are caused by the same linguistic mechanism(s). Here I will also argue that there are at least two behaviours that are prevalent amongst pejoratives but rare amongst laudatives. And, as I will argue in the last section, prohibitionism—of all the theories surveyed above—is the best equipped to explain this pattern in the data.

But before trying to rehabilitate prohibitionism, let’s first take stock of how laudatives and pejoratives are similar.

4 What laudatives do, too

4.1 Evaluating

For a start, the defining feature of both pejoratives and laudatives is that they can evaluate their targets. Like derogatory words, laudatory words

convey evaluative attitudes towards their targets. Unlike their derogatory counterparts, laudatives convey positive evaluations:

(6) Ann is *an angel*

In analogy to the notion of *derogatory force*, I suggest we call the positive-evaluative capacity of laudatives their *laudatory force*. Whereas derogatory force allows speakers to denigrate, disrespect and dehumanise a target of choice; laudatory force allows speakers to compliment, appreciate and admire.

There are more commonalities. For instance, derogatory force is typically *speaker-oriented* in the sense that the evaluation conveyed by pejoratives relays the speaker's evaluation of the target¹⁰. If Ann calls Bob a *dickhead*, she characterises him as an unsociable male and expresses her dislike of this characterisation. In other words, Bob is the target and Ann the owner of the negative evaluative attitude expressed by *dickhead*. Laudatory force is similarly speaker-oriented. Laudatives simultaneously describe and evaluate their targets. If Bob calls Ann a *sweetheart*, then he characterises her as pleasantly tempered and appreciates her for it. Here Ann is the target and Bob the owner of the positive evaluative attitude expressed by *sweetheart*.

In addition, laudatory words like derogatory words can convey their appraisal through expressing the speaker's emotions. Pejoratives can express sentiments such as disdain, anger, hatred and frustration. Calling someone an *asshole* or *dickhead* expresses disdain or contempt. Likewise, laudatives can express a slate of positively-valenced emotions including—but not limited to—appreciation, attraction, admiration, respect, infatuation and love. Calling someone a *sweetheart* or *darling* expresses affection. Calling someone a *genius* or *hero* expresses admiration.

4.2 Describing

Like pejoratives, some laudatives describe their targets, whereas other laudatives don't. Consider hybrid laudatives that attribute descriptive features to their targets while conveying positive evaluative attitudes. For instance, *stunner*, *stud* and *dreamboat* predicate and laud physical attractiveness.

There are also non-descriptive laudatives. Examples include interjections such as *yay*, *yippee*, *whoopee* and *wow*. On par with their derogatory counterparts, these interjections do not contribute any truth-evaluative content to the sentences they occur in. Instead, they

¹⁰E.g., see Potts (2007, pp.173–174) on speaker dependence.

signal the speaker’s positive appraisal of some contextually salient target. For instance, “I won the scholarship” and “Yippee, I won the scholarship” are logically equivalent.

Moreover, as is the case with hybrid pejoratives, speakers can felicitously use hybrid laudatives to either laud their intended targets without accurately describing them; or to accurately describe them without necessarily lauding them. For example, Bob may find Ann socially maladaptive and badly tempered, but may still felicitously call her a *sweetheart* or *darling* to express his affection towards her. Conversely, Ann may know Bob to be a physically attractive man, but may not appreciate him for this. Truthfully and unironically, she may describe Bob as a *looker* or *hunk* without expressing a laudatory attitude towards Bob’s good looks. She would not contradict herself were she to say: “Bob’s a hunk, but I don’t fancy him myself.”

4.3 Varying

So, like pejoratives, laudatives allow speakers to evaluate, emote and describe. Laudatory force further mimics several specific features of pejorative force. An oft-cited feature of pejorative force is its variability. Some words are ‘worse’ than others. [Hom \(2010\)](#) identifies three axes of variability along which, I think, laudatory force varies, too

Force variability To begin, pejorative force varies in its strength—that is, the degree to which a word negatively evaluates its target. This is best illustrated by co-referential pejoratives:

- (7) a. Ann is a *hag* < *bitch* < *cunt*
b. Bob is a *brute* < *bastard* < *asshole*

Laudatives are no different:

- (8) a. Ann is a *looker* < *babe*
b. Bob is a *whiz* < *mastermind*

This pattern is even more pronounced in laudatory adjectives:

- (9) That’s nice < great < amazing < phenomenal < sublime!

Historic variability A second axis of variability is historic. Pejoratives lose and gain derogatory force over time. The process by which linguistic items gain pejorative force is known as *pejoration*. Casually browsing the Oxford English Dictionary (OED) reveals that even the most forceful of words initially denoted perfectly inoffensive things. The n-word was derived from the Portuguese word for black. *Whore* shares its roots with the Latin for *dear*. Even *ass* comes from a Middle English nautical term, denoting “the bottom of a block”, having possibly alluded to the cleft through which a rope runs.

Pejoratives can lose their derogatory force through the process of *amelioration*. For example, members of younger generations will probably fail to recognise *queer* as a slur for gay men since it has undergone a mostly successful reclamation campaign by self-identified queers and well-meaning gender scholars. In contemporary parlance, *queer* is a relatively neutral label used to refer to anyone who isn’t cisgendered or heterosexual (Brontsema, 2014).

Theorists regularly include historic variability in their lists of explananda for explanatory adequate theories of pejoratives¹¹. The fact that laudatives undergo pejoration and amelioration has gone unmentioned. Though, again, casually browsing the OED proves illustrative. For instance, the exemplary laudatory adjective *nice* used to denote a slew of negative attributes. Here’s a small sample of its now obsolete senses:

Of a person: foolish, silly, simple; ignorant

Of an action, utterance, etc.: displaying foolishness or silliness;

Of conduct, behaviour, etc.: characterized by or encouraging wantonness or lasciviousness.

Other words have undergone amelioration in living memory. *Awesome*, *wicked*, *sick(ening)* and *insane* primarily had negative connotations several decades ago, but have acquired laudatory connotations as youth slang. Whereas a baby boomer might condemn a terrorist attack as *awesome*, a millennial would reserve the word for doling out praise.

With regards to the history of evaluative words, there is another notable parallel between laudatives and pejoratives. Both are produced by similar historical processes. It has been observed that many pejoratives acquire negative connotations as a result of being unflattering comparisons or metaphors. For example, Lepore & Stone

¹¹E.g., see Bolinger (2017), Brontsema (2014), Hom (2010), and Jeshion (2013).

(2018) single out unflattering metaphors as an important historical source of derogatory effects. They note that culturally prevalent negative attitudes towards a word's referent enter into our interpretations of its non-literal uses. For example, most of us agree that excrement is horrid. So, to call something *shit* is to call it horrid. Comparisons to genitalia and effluvia lend themselves well to conveying negative evaluations; the more taboo—the more negative the evaluation conveyed by the comparison. Instances of this phenomenon are self-explanatory, consider: *bitch*, *bastard*, *asshole*, and *motherfucker*.

Lepore and Stone only mention the interpretive effects of unflattering comparisons, but clearly flattering comparisons can convey evaluative attitudes, too. Comparisons to sweet, pleasant and divine things lend themselves well to conveying positive evaluations. Self-explanatory examples include *star*, *saint*, *angel* and *Adonis*. It thus seems that the processes by which pejoratives acquire derogatory force are sometimes responsible for the acquisition of laudatory force.

Moreover, the laudatives and pejoratives that result from these comparisons frequently retain both their literal and non-literal senses alongside each other. We are not wrong to call either Scylla or Heinrich Himmler *a monster*, even though only one of them is a frightening mythical creature. Nor are we wrong to call Brad Pitt a *stud*, even though he's not an uncastrated male horse kept for breeding. And both for laudatives and pejoratives, their evaluative non-literal readings sometimes eclipse its original non-evaluative literal meaning; as was the case with *nice*, which no longer means lasciviously dressed.

Valence variability The last axis of pejorative variability pertains to the direction of a pejorative's valence. Hom hints at two ways of interpreting valence variability. First, he observes that some pejoratives have an impressive affective range, being able to convey numerous emotions. Let's dub this feature a word's *affective variability*. Consider the broad affective range of pejorative interjections (Hom, 2010, p.165):

(10) *Damn!/Fuck!*

- a. I forgot my keys. (anger)
- b. I didn't get the job. (disappointment)
- c. That car is fast. (surprise)
- d. John is smart. (admiration)

Hom also observes that pejoratives can flip their valences from negative to positive (2010, p.165). Call this *valence variability*. This

occurs when pejoratives take on grammatical categories other than that of a noun, by doubling, for instance, as intensifiers to mean *really* or *very*:

- (11) a. Ann is a *damn* good/bad professor
b. Bob is a grown-*ass* man

Laudatives exhibit similar patterns of variability. For example, the laudatory exclamations appear to have the same expressive range as their pejorative analogues:

- (12) *Wow!/Oh my!*
a. I forgot my keys. (anger)
b. I didn't get the job. (disappointment)
c. That car is fast. (surprise)
d. John is smart. (admiration)

Some laudatory intensifiers can flip their valence from positive to negative as do pejorative intensifiers. Consider:

- (13) a. Ann is a *phenomenally* good/bad professor
b. Bob is *amazingly* funny/unfunny

This behaviour is restricted to a handful of grammatically flexible pejoratives and laudatives. Most pejoratives are inflexible, appearing only as noun phrases, not flipping their valences except for ironic or reclaimed uses. Similarly, laudatory noun phrases do not flip their valences either, again except for the occasional ironic use.

These observations so far clearly support [Nunberg's \(2018\)](#) and [Mišćević's \(2017\)](#) hypothesis that laudatives and pejoratives are mirror images of each other, different only in their direction of valence. In the remainder of this section, I argue that laudatives are not perfect mirror images of pejoratives. There are two behaviours that many pejoratives exhibit but that few laudatives emulate. But as I also argue below, these are merely differences in frequency of occurrence. These differences are therefore compatible with the hypothesis that pejoratives and laudatives very closely resemble each other.

4.4 Transgressing

The first *prima facie* difference between laudatives and pejoratives is that they are not equally socially transgressive. Laudatives, unlike pejoratives, are not typically subject to social norms that ban their use. Tellingly, most pejoratives get bleeped on daytime TV, whereas most laudatives do not. Swearwords and slurs, the subclasses of pejoratives whose use is socially transgressive by definition, appear to lack clear laudatory counterparts.

For a start, few laudatives classify as expletive. Only a handful of grammatically flexible expletives—words like *fuck*, *damn* and their derivatives—alternate between expressing derogatory and laudatory attitudes. Although exceedingly rare, some expletives exclusively convey positive attitudes. For example, American English *the shit*, British English *the dog's bollocks* and Iberian Spanish *de (la) puta madre* are expletive idioms that classify their targets as excellent or the best. The scarcity of expletive laudatives should be unsurprising. Laudatives rarely denote referents taboo enough to be considered expletive. That is, laudatives rarely pick out the aforementioned unmentionables such as bodily, sexual or blasphemous acts.

Laudatory analogues to slurs are equally sparse. Laudatives rarely, if ever, predicate social group memberships. Mišćević (2017, p.50) lists a few potential candidates: Croatian *Hrvatina*, American English *Yankee (Doodle)* and British English *John Bull*, which all roughly mean “a patriot admirable for their patriotism”. We could add *Aryan* to this list. When used by racists, *Aryan* can mean “a white person admirable for their racial purity”. We could also consider *lady* and *gentlemen* when used to laud someone for satisfying gendered expectations. *Lady*, for instance, can mean “a woman admirable for meeting cultural expectations of femininity”. Or, as Nunberg notes: some speakers appear to think that the word *lady* “demonstrates respect for the fair sex” (2018, p.274).

We could perhaps also include reclaimed slurs as instances of laudatory slurs such as the in-group uses of racist slurs by African Americans or of homophobic slurs by members of the LGBTQ+ community. These in-group uses can serve various non-derogatory purposes such as describing fellow group members, building rapport or signalling camaraderie in the face of shared experiences of animosity¹². However, it is unclear whether reclaimed slurs make convincing instances of laudatives. Reclaimed uses typically neutralise a slur's derogatory effects. In fact, successfully reclaimed slurs cease to be slurs,

¹²E.g., see Brontsema (2014) or Anderson (2018).

becoming neutral descriptive terms instead, as was the case with *queer*. In any case, laudatives that swear or slur are rare at best.

It would nonetheless be an oversimplification to conclude that laudatives, unlike pejoratives, are never subject to social prohibitions. Sometimes it is inappropriate to use laudatives. For example, it is uncontroversially inappropriate for professors to address their students as *babes* or *hotties*. This is because university employees are bound by social norms that forbid them from verbalising their physical attraction to students. Violations of these social norms can constitute sexual harassment.

Catcalling provides another example of prohibited laudatory language. Catcalling consists in expressing one's (sexual) attraction to strangers. Some catcalls are merely laudatory: a catcaller might address a passer-by as a *sweetie* or *babe*. Other catcalls are laudatory and lewd. Attitudes towards catcalling vary, but in some parts of the world catcalling is considered offensive. Various European countries have even criminalised it, fining catcallers hundreds of euros.

Laudatives that laud social group membership are similarly normatively suspect. For instance, hearers who reject white supremacy or rigid gender norms probably take offence at the use of *Aryan* or *lady* as laudatives. Dutch provides perhaps the clearest example of a ban on a slur-like laudative. The Dutch noun *blanke* roughly translates to *fair-skinned person* and histrionically derives from an evaluation-free proto-Germanic word, meaning *white* or *shiny*. In recent years, some Dutch speakers have begun rallying against its use, arguing that the word is objectionable for implying that fair-skinned people are racially superior. In other words, some Dutch speakers want to ban the word because they view it as a laudative that praises whiteness¹³.

Still, these examples are exceptions that prove the rule: laudatives, unlike pejoratives, rarely get banned.

4.5 Projecting

The second apparent difference is that laudatory force does not appear to project as readily as does derogatory force. The laudatory force of most laudatives *is* sensitive to the form of the sentence in which laudatives occur:

¹³I am indebted to Bart Geurts for suggesting this example to me.

- (14) a. **negation**: Ann is not a genius
 b. **conditional**: If Bob calls Ann daily, he’s a sweetheart
 c. **counterfactual**: Ann killed Snuggles but if she hadn’t, she would be an angel
 d. **tense**: Bob used to be a stud, but he has not aged well
- (15) a. **speech report**: Bob said “Ann is a cutie”
 b. **mental state report**: Ann thinks Bob is a sweetheart

The laudatives in these examples fail to heap praise upon anyone. The negation in (14a), for instance, prevents the laudative *genius* from lauding its intended target, Ann. The sentence also does not praise any other individuals who could be dubbed *geniuses*. The if-then construction in (14b) makes the expression of laudatory attitudes by the laudative *sweetheart* conditional on the sentence’s antecedent being true. As a result, the laudative does not express a laudatory attitude towards Bob or towards other potential sweethearts. In all these cases, the positive attitudes conveyed by the laudatives successfully embed under the operators in whose scope they appear.

Though, as I noted earlier, hearers sometimes take offence at the use of laudatives. If a laudative is offensive in virtue of clashing with social norms, it seems that it remains offensive in each instance of use. Consider expletive and slur-like laudatives:

- (16) a. **negation**: Ann’s not the dog’s bollocks
 b. **conditional**: If Ann obeys her man, then she is a lady
 c. **counterfactual**: Bob is half Jewish but if he had not been, he’d be Aryan
 d. **tense**: Bob’s parties used to be the shit

The expletive cases here clearly mirror the projective behaviours of their derogatory counterparts. Admittedly, the slur-like cases, *lady* and *Aryan*, are less convincing instances of projection. Hearers need to acknowledge various facts about the context of utterance, such as the speaker’s beliefs, to make the laudatives offensive. For example, *Aryan* is offensive only insofar hearers assume that speakers who use the word believe that being Aryan is being racially superior.

The Dutch slur-like laudative, *blanke*, provides a clearer case. Hearers who object to its use for praising whiteness tend to object to its use in all instances, irrespective of sentence form. In contrast to the above English-language examples, *blanke*’s laudatory attitudes appear to scope out from

embedding—lauding not only its intended targets, but anyone who fits the description of fair-skinned.

Projective behaviour is thus not as prevalent amongst laudatives as amongst pejoratives. However, it is not unattested either. And when laudatory terms do project, they do so for the same reasons as their derogatory analogues, namely for being expletive or otherwise normatively inappropriate.

4.6 Laudatives summarised

To summarise, laudatives and pejoratives closely resemble each other. Both laudatives and pejoratives can convey evaluative attitudes, affective states and descriptive contents. Both laudatives and pejoratives are speaker-oriented in the sense that they primarily enable speakers to express their attitudes. And both laudatives and pejoratives vary historically and comparatively with respect to the strength of the evaluations and emotions they express. Some laudatives can even flip their valence from positive to negative, just as some pejoratives can flip their valence from negative to positive.

It is clear that most laudatives most closely resemble non-slurring, non-expletive pejoratives such as *pencil pusher* and *freeloader*. Laudatives that denote taboo topics or social groups are extremely rare but not completely absent. As a result, there are few laudatory analogues to swearwords like *fuck* and *damn*, or to slurs like *kike* and *faggot*. Before concluding, let's consider the significance of these findings for the philosophical debate on derogatory language.

5 Theoretical consequences

5.1 A new desideratum for theories of pejoratives

For a start, these findings show that [Nunberg \(2018\)](#) and [Mišćević \(2017\)](#) were right to suggest that laudatives and pejoratives are inverse mirror images of each other, differing only in the direction of the evaluative attitudes they convey. It is therefore highly likely that the linguistic mechanisms that are responsible for communicating the derogatory attitudes of pejoratives are also responsible for communicating the laudatory attitudes of laudatives. As a consequence, theorists who wish to account for pejorative force should also account for laudatory force. Or, to rephrase this point as a demand: any empirically and explanatory adequate theory of pejorative force should

be able to explain—not only how pejoratives transmit negative attitudes—but also how laudatives transmit positive attitudes.

I think it will be relatively easy to retrofit existing theories of derogatory words to handle laudatory words. After all, the mechanisms that supposedly deliver derogatory force can in principle accommodate laudatory force. It is easy to imagine presuppositions, implicatures, tonal effects and expressive contents carrying positive instead of negative evaluative attitudes.

5.2 Consequences for prohibitionism

Of greater theoretical interest, I think, are the findings that (i) laudatives rarely ever project and (ii) that when they do, they seemingly do so because hearers consider their use normatively inappropriate. More specifically, my discussion of the projective behaviours of pejoratives and laudatives reveals a pattern in the projection data that could help theorists decide between the content-based and non-content-based approaches outlined in section 3.

The pattern is this: words that project are words that hearers, for whatever reason, judge to be normatively inappropriate. Hearers who object to the word *faggot* do so because its use contradicts their normative belief that speakers ought never to derogate individuals for being gay. And hearers who object to the use of the words *Aryan* or *blanke* do so because using these words violates the hearers' normative belief that speakers ought never to laud individuals for being white. Likewise, hearers who object to the use of *dickhead* or *the shit* do so because uttering these words violates social norms that ban speakers from mentioning the 'unmentionables' *dick* and *shit* in polite company. At the same time, words that don't project—be they laudatives or pejoratives—are not socially transgressive. Exemplary here are the words *genius*, *angel*, *jerk* and *brute*. So, only normatively inappropriate words appear to project.

This pattern, I think, should be cause for celebration amongst prohibitionists and cause for concern amongst content-theorists. The pattern neatly fits with prohibitionism for a clear reason: because prohibitionism provides a plausible and parsimonious explanation of how it arises. That is, according to prohibitionists, a word projects insofar it offends the hearer and a word offends insofar the hearer deems the word to be normatively inappropriate or forbidden. In fact, prohibitionism appears to be the only theory that correctly predicts this pattern since it is the only theory that singles out prohibitions as indicative of offensiveness.

5.3 Consequences for content-based theories

By contrast, the pattern cannot be meaningfully explained in terms of a word's linguistic content alone. This is because hearers' judgements about the normative appropriateness of words clearly go beyond linguistic content. Or, to put it differently: a hearer's social norms are not encoded into the linguistic meaning of words. For instance, the linguistic meaning of homophobic slurs does not encode that speakers are banned from expressing negative attitudes towards gay people. Nor do the contents of expletives like *dickhead* or *the shit* prohibit speakers from expressing their attitudes through comparisons to genitalia or excrement. To me this suggests that linguistic content alone doesn't suffice to explain why some pejoratives and laudatives project, whereas others don't.

At least two prominent advocates of content-based approaches appear to agree with me on this point. I say this because I think that a common content-theoretic account to projection implicitly endorses a minimal form of prohibitionism. That is, the content-theorists, [Camp \(2013\)](#) and [Hom \(2020\)](#), share the intuition that projection occurs when hearers are confronted with the disturbing attitudes encoded into the linguistic contents of pejoratives. On this view, projection reflects the psychological distress hearers experience when they encounter the disturbing contents of a pejorative.

However, to accept that projection consists in hearers taking offence at derogatory content is to accept a form of prohibitionism. This is because prohibitionism, on my view, is the best theory we have of what it means to consider a word offensive, disturbing, transgressive or otherwise objectionable. To reiterate: according to prohibitionists, words offend insofar they are inappropriate or forbidden. Hom seems to share this prohibitionist sensibility: he specifically invokes the normative inappropriateness of slurs to explain why hearers deem them offensive in all instances of use. Slurs, he writes, offend hearers because they “(1) force hearers to entertain a degenerate way of classifying the world, and (2) they signal that they themselves approve of this classificatory scheme as *normatively appropriate*” (2020, p.13, emphasis added). It thus seems that accounting for the distribution of projective behaviours amongst evaluative words forces even content-theorists like Hom to accept a minimal form of prohibitionism.

To be clear, by *minimal prohibitionism* I mean the position that hearers' beliefs about which words are (in)appropriate or forbidden determine which words offend hearers. This view is not to be confused with *full-blown prohibitionism*—the position that prohibitions on pejoratives fully explain why pejoratives convey derogatory attitudes.

There are good reasons to reject full-blown prohibitionism. Even [Lepore & Stone \(2018\)](#), prohibitionism’s most ardent supports, don’t ascribe to it; since, as [Camp \(2013\)](#) stressed, it can’t explain why we often ban words *because* they convey attitudes that we find objectionable. In other words, a good theory of evaluative words should explain how pejoratives and laudatives come to carry evaluative attitudes to begin with. For this reason, Lepore and Stone prop up prohibitionism with their tonal theory—their tonal theory explains how words come to convey evaluative attitudes; their prohibitionism explains how some words come to offend.

Minimal prohibitionism thus differs from full-blown prohibitionism in that it remains neutral about the mechanisms by which words convey evaluative attitudes. It merely tells us what it means for us to be offended by a word. Minimal prohibitionism therefore is compatible with approaches that encode evaluative attitudes into the linguistic contents of evaluative words. And this minimal prohibitionism, I believe, is needed to explain why some evaluative words project and why others don’t.

6 Conclusion

In conclusion, my comparison of laudatives and pejoratives has demonstrated that [Nunberg’s \(2018\)](#) and [Mišćević’s \(2017\)](#) hypothesis is correct: laudatives do really mirror pejoratives. I have also demonstrated that their mirroring is not perfect. Although many pejoratives offend their hearers and project their evaluative attitudes, few laudatives do. I have argue that these minor dissimilarities are best explained in terms of non-content-based theories in lieu of content-based ones. However, the main takeaway here is not that prohibitionism is the best theory we have of projection. Rather, the main contribution of this paper is the observation that the projective behaviours of all evaluative words—including laudatives—can inform and contain theorising about derogatory language.

Another takeaway message is that theorists are in need of a better more in-depth analysis of the projective behaviours of evaluative words. I have attempted to disentangle what projects out when pejoratives and laudatives project, but numerous unanswered questions remain—amongst them: How do the projective behaviours of expletives and slurs differ? Is the projection of laudatory attitudes more common in other languages? Could there be any other elements of the meaning of evaluative words that escape embedding? If so, which ones?

References

- Anderson, L. (2018). Calling, addressing, and appropriation. In D. Sosa (Ed.), *Bad words: Philosophical perspectives on slurs* (pp. 6–28). Oxford University Press. doi: 10.1093/oso/9780198758655.003.0002
- Anderson, L., & Lepore, E. (2013). What did you call me? Slurs as prohibited words. *Analytic Philosophy*, 54(3), 350–363. doi: 10.1111/phib.12023
- Blakemore, D. (2015). Slurs and expletives: a case against a general account of expressive meaning. *Language Sciences*, 52(1), 22–35. doi: 10.1016/j.langsci.2014.06.018
- Bolinger, R. J. (2017). The pragmatics of slurs. *Nous*, 51(3), 439–462. doi: 10.1111/nous.12090
- Brontsema, R. (2014). A queer revolution: Reconceptualizing the debate over linguistic reclamation. *Colorado Research in Linguistics*, 17(1), 1–17. doi: 10.25810/dky3-zq57
- Camp, E. (2013). Slurring perspectives. *Analytic Philosophy*, 54(3), 330–349. doi: 10.1111/phib.12022
- Cepollaro, B. (2015). In defence of a presuppositional account of slurs. *Language Sciences*, 52(1), 36–45. doi: 10.1016/j.langsci.2014.11.004
- Croom, A. M. (2013). How to do things with slurs: Studies in the way of derogatory words. *Language & Communication*, 33(3), 177–204. doi: 10.1016/j.langcom.2013.03.008
- Croom, A. M. (2014). The semantics of slurs: a refutation of pure expressivism. *Language Sciences*, 41, 227–242. doi: 10.1016/j.langsci.2013.07.003
- Geurts, B. (2007). Really fucking brilliant. *Theoretical Linguistics*, 33(2), 209–214. doi: 10.1515/tl.2007.013
- Hom, C. (2008). The semantics of racial epithets. *Journal of Philosophy*, 105(8), 416–440. doi: 10.5840/jphil2008105834
- Hom, C. (2010). Pejoratives. *Philosophy Compass*, 5(2), 164–185. doi: 10.1111/j.1747-9991.2009.00274.x
- Hom, C. (2012). A puzzle about pejoratives. *Philosophical Studies*, 159(3), 383–405. doi: 10.1007/s11098-011-9749-7

- Hom, C. (2020). Slurs, assertion, and predication. In S. Goldberg (Ed.), *The Oxford handbook of assertion* (pp. 286–305). Oxford University Press. doi: 10.1093/oxfordhb/9780190675233.013.32
- Hom, C., & May, R. (2013). Moral and semantic innocence. *Analytic Philosophy*, 54(3), 293–313. doi: 10.1111/phib.12020
- Jeshion, R. (2013). Expressivism and the offensiveness of slurs. *Philosophical Perspectives*, 27(1), 231–259. doi: 10.1111/phpe.12027
- Kaplan, D. (1999). *What is meaning? Explorations in the theory of meaning as use* (Unpublished master’s thesis). UCLA.
- Lepore, E., & Stone, M. (2018). Pejorative tone. In D. Sosa (Ed.), *Bad words: Philosophical perspectives on slurs* (pp. 132–154). Oxford University Press. doi: 10.1093/oso/9780198758655.003.0007
- Mišćević, N. (2017). The semantic structure of pejoratives. *Linguistik Online*, 82(3), 49–64. doi: 10.13092/lo.82.3714
- Nunberg, G. (2018). The social life of slurs. In D. Fogal, D. W. Harris, & M. Moss (Eds.), *New work on speech acts* (pp. 238–291). Oxford University Press. doi: 10.1093/oso/9780198738831.003.0010
- Popa-Wyatt, M. (2016). Not all slurs are created equally. *Phenomenology and Mind*, 11(1), 150–156. doi: 10.13128/Phe_Mi-20115Web
- Potts, C. (2007). The expressive dimension. *Theoretical Linguistics*, 33(2), 165–198. doi: 10.1515/tl.2007.011
- Pullum, G. K. (2018). Slurs and obscenities: Lexicography, semantics and philosophy. In D. Sosa (Ed.), *Bad words: Philosophical perspectives on slurs* (pp. 165–192). Oxford University Press. doi: 10.1093/oso/9780198758655.003.0009
- Rappaport, J. (2019). Communicating with slurs. *The Philosophical Quarterly*, 69(277), 795–816. doi: 10.1093/pq/pqz022
- Richard, M. (2008). *When truth gives out*. Harvard University Press. doi: 10.1093/acprof:oso/9780199239955.001.0001
- Schlenker, P. (2007). Expressive presuppositions. *Theoretical Linguistics*, 33(2), 237–45. doi: 10.1515/tl.2007.017
- Sennet, A., & Copp, D. (2014). What kind of a mistake is it to use a slur? *Philosophical Studies*, 172(4), 1079–1104. doi: 10.1007/s11098-014-0338-4

Stapleton, K. (2010). Swearing. In M. A. Locher & S. L. Graham (Eds.), *Interpersonal pragmatics (Handbooks of Pragmatics)* (pp. 289–306). De Gruyter Mouton. doi: 10.1515/9783110214338.2.289