Three Rich-Lexicon Theories of Slurs: A Comparison

Abstract: Many authors writing on slurs think that they are lexically rich, in the sense that their lexical meaning comprises both a descriptive dimension and an expressive/evaluative one, the latter accounting for their derogatory character. However, more fine-grained theories of slurs have recently been proposed, drawing on frameworks from lexical semantics. My main aim in this paper is to compare three such fine-grained rich-lexicon theories – the one put forward by myself in previous work with two similar ones, Croom's (2011, 2013) and Neufeld's (2019, 2022). While my primary aim is to underlie the commonalities and differences between the three views, I also aim at showing that the latter views have some issues that are easily avoided in the framework I propose.

Keywords: slurs, lexical semantics, rich-lexicon theories, derogatory and non-derogatory uses of slurs

Many authors writing on slurs think that they are lexically rich, in the sense that their literal meaning comprises at least a descriptive aspect and an expressive or evaluative one, the latter accounting for their derogatory character. According to such views, a slur like "boche", for example, refers to Germans (i.e., its descriptive part is the class of German people), while the expressive or evaluative part contains the negative attitudes or evaluations the speaker harbours against Germans. The second aspect can be encoded in the lexical meaning of the slur either as a semantic component (for example, in views that postulate a single layer of meaning, perhaps composed of a conjunction of both descriptive and expressive/evaluative elements, or more than one layer of meaning), or as a pragmatic component (for example, in views that take derogation to be explained by the slur triggering certain presuppositions, licensing certain implicatures/inferences, residing in the speech acts made, etc.).

Recently, however, more fine-grained theories of slurs have been proposed, drawing on frameworks from, or similar to those from, *lexical semantics*. Lexical semantics is the branch of semantics that studies the structure of the meaning of words. Two broad approaches have been prominent in this field: rich-lexicon ones and thin-lexicon ones. According to the former, the lexical entry of a certain type of word is complex, consisting in various meaning elements connected in certain ways, among which one or several are selected as the word's sense in a certain context; according to the latter, the lexical entry of a certain type of word is thin – at limit consisting only in a syntactic string – and gets enriched by further meaning elements in a certain context, based on encyclopaedic knowledge and contextual clues. While the debate

about which of these two broad approaches is correct and about which of the many theories on the market best captures the phenomena is in full swing, the application of any of these to slurs is relatively new. Thus, my main purpose in this paper is to compare three rich-lexicon theories of slurs that have surfaced in the literature recently: those by Croom (2011, 2013, 2014, 2015, 2018), Neufeld (2019, 2022), and Zeman (2022), with the aim of underlying their commonalities and differences, and of showing that the two former views have some issues that are easily avoided in that proposed in Zeman (2022). Discussing these theories under the same umbrella further shows the fruitfulness of the rich-lexicon approach – already a successful endeavour in relation to many natural language expressions and in tackling various issues and phenomena such as word meaning, context-sensitivity, polysemy and co-predication, metaphor etc. – and puts the basis of a research program that approaches slurs by investigating basic aspects of their meaning and the interconnections between them. As a byproduct of the comparison undertaken, several aspects of my proposal in Zeman (2022) will be further clarified and modified.

The paper unfolds as follows. In section 1, I give a quick characterization of slurs and of their most common uses, which will constitute the background against which the three theories mentioned will be compared. In section 2, I spell out what a rich-lexicon theory comes down to, using nouns as illustration. In the following section, I show one way in which such a theory can be applied to slurs, using my own previous work (Zeman (2022)). In section 4 and 5, I present Croom's and Neufeld's views, and raise a few objections to each of them. I conclude, in a short section 6, by comparing the three rich-lexicon approaches scrutinized.

1. Slurs and slurring

Slurring is widespread in many of the world's cultures, and slurs exist in most languages (including sign languages). It is agreed in the literature that the function of slurs is "to derogate or dehumanize (…), to signal that their targets are unworthy of equal standing or full respect (…), that they are inferior as persons" (Jeshion, 2013a: 232). Slurs target both groups and individuals – the latter via group membership, and are based on a plethora of extant and perceived characteristics of the members of the target group which are evaluated negatively – such as race, sexual orientation, gender, nationality, ethnicity, disability, political views, character, personality, physical traits, psychological traits, and so on.

While the main function of slurs is to derogate, there are other uses of slurs, many non-derogatory. Essentially, authors have distinguished two broad categories of uses of slurs: derogatory, also known as "weapon uses", and non-derogatory uses, also known as "non-

weapon uses" (Jeshion (2013a). The standard derogatory use is usually exemplified with sentences that simply predicate the slur of a person, such as "Hans is a boche."[1] Or "You are a boche.", or in which someone is called using a slur, as in "You, boche!". There are also what could be called *internalized* derogatory uses, in which a person belonging to the target group uses the slur to self-derogate or to derogate other members of the group: e.g., "I'm a boche.".

When it comes to non-derogatory uses, there are several types that have been acknowledged in the literature. Among these, the most studied (see Saka (2007), Hom (2008, 2010), Croom (2011, 2013, 2014, 2015, 2018), Anderson & Lepore (2013a), Jeshion (2013a, 2020), DiFranco (2014), Bianchi (2014), Ritchie (2017), Cepollaro (2017b, 2020), Anderson (2018), Burnett (2020), Popa-Wyatt (2020), Cervone, Augoustinos & Maass (2021), Jusińska (2021) – among many others) are *appropriated* uses (or reappropriated, or reclaimed – the terminology varies). "Queer" stands out as the main example of a slur that has been successfully appropriated, as it seems to have entered common academic usage (there is a journal titled *Queer Studies*), entertainment (there is a TV series called *Queer Eye*), and daily non-derogatory use. The n-word is also frequently given as an example of successful appropriated use, although in this case its success is measured by it being taken up within the African American community only, and not at large. Appropriation is itself a complex, multifaceted phenomenon, and while there are many ways to use slurs appropriately (see, for example, the rich essays by Naylor (1986), Kennedy (2003) or Brontsema (2014)), two subtypes have been singled out by philosophers: i) as used in political activism, as a tool against oppression and discrimination; ii) as terms of endearment, signalling familiarity and camaraderie. Besides appropriated uses, other non-derogatory uses of slurs are what Hom (2008) has called "non-derogatory non-appropriated", or corrective uses, exemplified by sentences like "Institutions that treat Chinese as chinks are morally depraved." (Hom, 2008: 423); metaphorical uses, exemplified by the title of the famous song by John Lennon & Yoko Ono, "Woman is the n*** of the world."; didactical/quotational uses, such as these employed in academic works, courts of law, etc.

There is also another subclass of non-derogatory uses of slurs that is somewhat different than the ones above. Anderson (2018) has noted, following work by Smitherman (2006), that the n-word (in its non-rhotic form) can be used in certain contexts *referentially* to mean "friend", "buddy", "man" (Anderson, 2018: 9). Thus, Anderson writes: "the use[s] of the [n-word] commonly attributed to rappers, black comedians, and black youth (…) aren't employed

---

[1] As customary, I will try to keep the mention of slurs at a minimum. I'm using "boche" here as an example due its reduced (if any) derogatory power.

non-referentially to illuminate claims about the persistence and reality of racism, but are often used referentially in a friendly manner." (Anderson 2018: 7). Although Anderson claims that such uses are "not typically interpreted as one of derogation, but instead as one of camaraderie" (Anderson, 2018: 7), thus likening them to appropriated uses, he also claims that "there are a variety of uses black speakers employ with [the n-word, in its non-rhotic form], some of them positive, some neutral, and others negative" (Anderson, 2008: 14). This neutral use of the n-word corresponds to one of the uses of it distinguished by Smitherman and characterized as "generic, neutral reference to African Americans" (Smitherman, 2006: 52). In Zeman (2022), I pointed to a yet different type of use, based on data about the use of ethnic slurs like "țigan" (roughly translated in English as "gipsy") in Romanian and in other languages from Eastern Europe by members of the Roma community to identify as belonging to a certain ethnic group. The data was gathered from a Romanian Government Report (2009) (henceforth RGR) and consists in answers to questions Roma ethics were asked, such as the following: "How do you usually refer to yourself and to those in your community?", with the corresponding percentage of answers: "țigan" – 66%; "rom" (the word for "gypsy" in the Roma ethnics' language, Romani) – 30%; no answer – 4% (RGR: 46); "How do you refer to yourself and to those in your community in the Romani language?": "țigan" – 44%; "rom" – 32%; no Romani speaker – 22%; no answer – 2% (RGR: 46); "How offensive do you find the word "țigan"?", on a scale from 1 ("not offensive at all") to 5 ("very offensive"): 1 – 38.3%; 5 – 26% (RGR: 27). The data is thus robust, and supports the postulation of a sui-generis, different type of use of slurs, which Zeman (2022) has dubbed *identificatory*.[2]

Importantly for us, the difference between these two types of uses and the ones showcased above is that the former don't aim at changing the negative evaluation encoded in the slur into a positive one. While appropriated (and perhaps corrective and metaphorical uses) aim to put a positive spin on a slur, referential and identificatory uses are simply neutral and

---

[2] The two types of uses presented in the paragraph seem quite close, so an interesting question is whether they are, after all, different. In Zeman (2022), I argue that identificatory uses are more focused on identifying the group those involved belong to (along an objective dimension – say, race or ethnicity) than referential uses (even though, of course, the latter also serve to single out a particular group). In other words, while the main purpose of referential uses is to select a group of people the speaker is close to, that of identificatory ones is to select those sharing a certain objective trait – e.g., ethnicity. One further consideration that supports distinguishing these two uses is that with referential uses, but *not* with identificatory ones, one can refer to someone that is not part of the group. In what follows, I will assume that referential and identificatory uses of slurs are two different types of uses of slurs. If one doesn't agree with this claim, then the separate treatment given in section 3 to these two types of uses will collapse into the one proposed for identificatory uses. See Zeman (2022) also for arguments that this type of use cannot be reduced to the other ones acknowledged in the literature or that it merely amounts to a conversational shortcut.

their purpose it to merely refer to members of the target group, either as individuals/fellow humans or as belonging to a certain social group.

Theories of slurs can be and are often grouped in two broad categories: semantic theories and pragmatic ones. The former (Richard (2008), Hom (2008, 2010, 2012), Hom and May (2013, 2018), Predelli (2013), Camp (2013, 2018), Jeshion (2013a, 2013b, 2016, 2018), Vallée (2014), Scott and Stevens (2019), Orlando and Saab (2020), Sullivan (2022), etc.) situate the derogatory element in the linguistic/literal meaning of the slur; the latter (Potts (2005, 2007), Schlenker (2007), Whiting (2007, 2013), Williamson (2009), McCready (2010), Gutzmann (2015, 2019), Cepollaro (2015, 2020), Cepollaro and Stojanovic (2016), Nunberg (2018), Marques and García-Carpintero (2020), etc.) situate the derogatory element at the level of pragmatics, as conveyed via presuppositions or various kinds of implicatures. Many other pragmatic theories appeal to mechanisms other than the classical ones. For example, a group of views takes the derogatory content to be arrived at by common inferences drawn within the communities using the term (Tirrell (1999), Hornsby (2001)). Other views take it to reside in the "social meaning" slurs have (Burnett (2020)), leading to a focus on the social mechanisms by which derogation is realized and propagated (e.g., Popa-Wyatt and Wyatt (2018), Hess (2020)). Bolinger (2017) argues for a view according to which derogation is a matter of choosing between alternative linguistic means to express oneself. Another type of pragmatic views focuses on the type of speech acts speakers perform in uttering slurs (Langton (2012), Bianchi (2018), Kirk-Giannini (2019), Liu (2021)), while yet another appeals to the notion of "register" (Diaz Legaspe, Liu and Stainton (2020)). Finally, there are also more radical pragmatic views, such as that of Anderson and Lepore (2013a; 2013b), who take slurs to be taboo words, situating their derogatory character in the societal configurations that prohibit the use of certain words, and that of Stojnić and Lepore (2022), who explain the derogatory effect of slurs by the associations articulations of slurs elicit both in their users and in members of the target group.

The type of theories I will focus on in this paper – lexical theories – can be said to incorporate both semantic and pragmatic elements and thus fall in between the two categories I started this classification with. Lexical theories generally distinguish between the *lexical entry* of a word (or its lexical, or literal, meaning) and its *sense* (the meaning it has on a particular occasion of use, in a certain context). Regardless of what the postulated relation between these is, lexical theories could be interpreted both as semantic (the lexical meaning of a word *is* its semantic meaning, which is the meaning grasped by speakers of a language when they know a word), even if semantics here is not conceived along familiar, truth-conditional lines, and as

pragmatic (a word can have different senses in different contexts) – that is, pertaining to a word's use. It might be that, according to more exact definitions of "semantic" and "pragmatic", the type of theories I will discuss below will neatly fall in one camp or another. However, it is not my aim in this paper to settle this vexed issue; its purpose is rather to give a clear picture of the type of theory I focus on and its application to slurs – regardless of how its instantiations end up being classified.

2. What Is a Rich-lexicon Theory?

In lexical semantics, there is currently a debate opposing "rich-lexicon" theories to "thin-lexicon" ones (for foundational work, see Jackendoff (1990), Pustejovsky (1995), Asher (2011); for review articles and arguments for one type of theory or the other, see Falkum and Vicente (2015, 2020), Vicente and Falkum (2017), Ortega Andrés and Vicente (2019), Hogeweg and Vicente (2020), Vicente (2010, 2018, 2021)). Most of the debate is concerned with polysemy, a phenomenon the two broad categories of views mentioned aim to give the most satisfactory account of. Although I think that slurs and other evaluative expressions are polysemous, I won't defend this claim here; nor will I be concerned more generally with this particular phenomenon. My focus in this section is to present the main characteristics of the family of views belonging to the rich-lexicon family.

What is a rich-lexicon theory, then? A common claim made by proponents of such theories is that a word's lexical entry comprises various dimensions of meaning, among which the one that an expression has in a certain context (its "sense") is selected. In contrast, a common claim made by the proponents of thin-lexicon theories is that a word's lexical entry is basic, with the sense of a word in a certain context being constructed out of that basic meaning by using encyclopaedic knowledge, contextual clues, etc. or by other ways of enriching it. How to determine the meaning dimensions that constitute a word's lexical entry according to rich-lexicon theories is a crucial issue. In the case of nouns, one common way to do so has been to appeal to the old Aristotelian idea of *qualia*: units of meaning that encapsulate certain types of information and which are interrelated in certain ways (e.g., Pustejovsky (1995)). Thus, various rich-lexicon theorists have taken the lexical entry of nouns to comprise at least the following: perceptual information about the objects referred to with the noun (I will refer to it with the label PERCEPTUAL), information about what those objects are made of or their parts (CONSTITUTIVE), how they came to being or the purpose of their creation (AGENTIVE), their typical function (TELIC), etc. For purposes of illustration, I will use the framework proposed by Del Pinal (2018), who takes the meaning of nouns to be tuples of extensions ("E-

structures") and conceptual structures ("C-structures"). In that framework, and to give a common example, the lexical entry for the noun "book" will be rendered as

$$
\begin{bmatrix}
\text{PERCEPTUAL: shape, size, color etc.} \\
\text{CONSTITUTIVE: cover, pages, etc.} \\
\text{AGENTIVE: printed, written etc.} \\
\text{TELIC: to entertain, to learn from etc.}
\end{bmatrix}
$$

While this structure represents the full literal meaning of "book", various senses of the word are selected in different contexts. For example, in a context c in which "book" is part of the sentence "This book is bound in leather.", the CONSTITUTIVE dimension is selected; clearly, what is relevant here is what the book is made of. This is represented in the framework adopted as

$$
\begin{bmatrix}
\text{PERCEPTUAL: shape, size, color etc.} \\
\textbf{CONSTITUTIVE: cover, pages, etc.} \\
\text{AGENTIVE: printed, written etc.} \\
\text{TELIC: to entertain, to learn from etc.,}
\end{bmatrix}^{c}
$$

with the bolded dimension of meaning signifying that it has been selected in context c. In contrast, in a context c' in which "book" appears in the sentence "This book has influenced a whole generation of thinkers.", what is relevant is the book's content, and so the TELIC dimension is selected. This is represented in the framework adopted as

$$
\begin{bmatrix}
\text{PERCEPTUAL: shape, size, color etc.} \\
\text{CONSTITUTIVE: cover, pages, etc.} \\
\text{AGENTIVE: printed, written etc.} \\
\textbf{TELIC: to entertain, to learn from etc.}
\end{bmatrix}^{c'}
$$

In the examples above, linguistic material from the two sentences "book" appears in ("bound in leather" and "has influenced a whole generation of thinkers", respectively) gives an indication of which sense of the word is the most relevant in that context. However, "book" can appear in sentences in which the surrounding linguistic material is not as generous with the clues (an example would be "This book is good." – ignore for the moment the various senses

"good" can have). In cases in which linguistic clues are missing, senses are selected via relevance or other contextual cues (see, among others, Carston (2021)).

What are the mechanisms by which such selection takes place is, obviously, another crucial matter. A variety of proposals can be found in the literature: see, for example, those advocated or discussed in Pustejovsky (1995), Frisson (2009), Asher (2011), Schumacher (2013), Del Pinal (2018), etc. I won't go into this issue here, despite its important place in any rich-lexicon theory. Instead, I will refer to the mechanism(s) by which the meaning dimensions in the lexical entry of a word are selected as the word's sense in a given context with the placeholders "foregrounding" and "backgrounding". Using this terminology, I will say that the CONSTITUTIVE dimension of "book" is foregrounded when used in the sentence "This book is bound in leather." in *c*, while the other meaning dimensions are backgrounded, and that the TELIC dimensions of "book" is foregrounded when used in the sentence "This book has influenced a whole generation of thinkers." in *c'*, while the other meaning dimensions are backgrounded.[3]

It is worthwhile noting that this approach to nouns offers answers to a host of questions found at the intersection of philosophy, linguistics and cognitive psychology (for a survey, see Hogeweg & Vicente (2020)). Thus, postulating rich lexical meanings gives a straightforward answer to the question of what word meanings are: they are rich configurations of basic elements that interconnect in principled ways. This approach also offers a treatment of various phenomena such as coercion effects (meaning shifts as those in sentences like "The ham sandwich wants the bill."; logical metonymy as in sentence like "Sarah began the book."), non-homogenous predication (some predicates attach to *aspects* of an object, not to the entire object), multi-dimensional modification (as with "stone lion"), radical truth-conditional variation (as in Travis cases – "green leaves"), etc. It has also been considered one of the main ways to explain polysemy (as we saw with "book") and co-predication (see footnote 3). In all these cases, the mechanisms of foregrounding and backgrounding illustrated above play a major role. While the current theories postulating rich meanings are not without problems (Hogeweg & Vicente's (2020) conclusion in the survey mentioned is rather cautious), this approach has been one of the main contenders in the debate over the issues mentioned, with

---

[3] It is also possible that two meaning dimensions are simultaneously selected in a given context: for example, in the sentence "This beautifully bound book has influenced a whole generation of thinkers.", both the CONSTITUTIVE and the TELIC dimensions are foregrounded. (The phenomenon illustrated by this sentence is an example of co-predication (see below), and the literature on it substantial; for a recent discussion, see Ortega Andrés and Vicente (2019)). It is thus not a claim of a rich-lexicon theory that in a given context at most one meaning dimension is selected as the word's sense.

significant success. Extending it to expressions it hasn't been so far applied to in order to solve similar problems would augment the framework's reach and consolidate its appeal. Seen from the other side, treating slurs as incorporating rich meanings makes room for a novel approach to this type of expression that aims to solve the particular issues that they pose (accounting for slurs' various uses, or for changes in their meaning, etc.) and neatly connects them with other natural language expressions that have received a rich-lexicon treatment. With this idea of the fruitfulness of the rich-lexicon framework in mind, and with the barebones of its main characteristics on the table, we can move now to its application to slurs.

3. Applying the rich-lexicon framework to slurs

The application of the model of word meaning presented above to slurs proceeds by introducing in the conceptual structure associated with them (del Pinal's "C-structure") an evaluative dimension (EVALUATIVE) which is responsible for the slurs' derogatory character. What exactly enters that dimension is, obviously, an important issue and up for debate, but – again – I won't settle it here. For example, I will remain neutral between a purely *expressive* approach, according to which the dimension should comprise merely attitudinal elements, such as speakers' attitudes towards the target group (like contempt) and a purely *informational* one, according to which the dimension should comprise information such as "detestable because of being X": for the purposes of this paper, any of the two would do. Rather, my aim is to show how, starting from the rich lexical entries of nouns espoused above, one can arrive, with some modifications, at rich lexical entries suitable for slurs (since slurs are nouns).

Thus, in consonance with the framework described in the previous section, I take the lexical meaning of a slur to be the rich lexical configuration that comprises all its meaning dimensions. Here is an attempt at illustrating what the general structure of the lexical meaning of a slur would come down to in this framework:

EVALUATIVE: negative evaluation of the members of the target group.[4]

PERCEPTUAL: shape, size, color etc.

CONSTITUTIVE: body parts etc.

AGENTIVE: born etc.

ORIGIN: provenance, history, social standing etc.

---

[4] Here is one point in which the view differs from the one in Zeman (2022), where EVALUATIVE also comprised information about the stereotypes associated with the target group, about the ideology of the users of the slur, etc.

The novelty is, of course, the introduction of the EVALUATIVE dimension, but also the replacement of the TELIC one from the lexical meaning of the nouns considered above (since it seems meaningless to ask what the telos of a social group, to which slurs refer to, is) with ORIGIN, a dimension comprising information about the provenance, history and social standing of the target group. Using the notions of foregrounding and backgrounding introduced in the previous section, one can see how both derogatory and neutral uses of slurs are represented in the framework adopted, as well as the main difference between them. Thus, a derogatory use of a slur in a context c will be represented as follows, with the bolded dimension of meaning signifying that it has been foregrounded in that context, while the non-bolded dimensions of meaning signifying that they have been backgrounded in that context:

$$
\begin{bmatrix}
\textbf{EVALUATIVE: negative evaluation of the members of the target group.} \\
\text{PERCEPTUAL: shape, size, color etc.} \\
\text{CONSTITUTIVE: body parts etc.} \\
\text{AGENTIVE: born etc.} \\
\text{ORIGIN: provenance, history, social standing etc.}
\end{bmatrix}^{c}
$$

A neutral use of the same slur – for example, an identificatory one – in a different context c' will be represented as follows, with the bolded and non-bolded dimensions of meaning having the same significance as above:

$$
\begin{bmatrix}
\text{EVALUATIVE: negative evaluation of the members of the target group.} \\
\text{PERCEPTUAL: shape, size, color etc.} \\
\text{CONSTITUTIVE: body parts etc.} \\
\text{AGENTIVE: born etc.} \\
\textbf{ORIGIN: provenance, history, social standing etc.}
\end{bmatrix}^{c'}
$$

The difference between these two uses is captured in the theory by the foregrounding of different meaning dimensions: in derogatory uses, EVALUATIVE is foregrounded, and the other dimensions backgrounded, showing that the speaker's negative evaluation of the members of the target group is what they intend to communicate, while in identificatory uses ORIGIN is foregrounded and the others (importantly, EVALUATIVE as well!) backgrounded,

showing that the speaker's evaluation of the members of the target group is not part of what they intend to communicate.[5]

Now, an interesting question that arises is how appropriated uses are treated in this framework. Recall that in appropriated uses, in contrast to neutral ones, the aim is to re-signify the slur by transforming it to encode a positive evaluation instead of a negative one. So, as in derogatory uses of a slur, an evaluation is present in appropriated ones too, which is expressed in the theory by foregrounding EVALUATIVE when dealing with the latter. The change in evaluation is captured by appeal to the notion of *valence* (used by many authors – see, for example, Cepollaro (2017a), Cepollaro and Stojanovic (2016), etc.), which is construed as a function from a context c to the type of evaluation (positive or negative) relevant in c. Thus, the correct rendering of EVALUATIVE will be something along the lines of "valence (c) = positive/negative evaluation of the members of the target group", with the appropriated use of a slur in a context c'' being represented as follows:

c''

EVALUATIVE: positive evaluation of the members of the target group.
PERCEPTUAL: shape, size, color etc.
CONSTITUTIVE: body parts etc.
AGENTIVE: born etc.
ORIGIN: provenance, history, social standing etc

Importantly, and to compare how the uses mentioned so far are represented, while in identificatory uses of slurs EVALUATIVE is backgrounded, both in derogatory and in appropriated uses EVALUATIVE is foregrounded. Finally, it is important to note that the framework can easily account for other types of uses of slurs from among those mentioned in section 2. In metaphorical uses of slurs, for example, ORIGIN is backgrounded, which allows application of the slur to individuals that do not belong to the target group, while the other dimensions (including EVALUATIVE) can be foregrounded – which exactly depending on the aspect of the slur's meaning that is used to base the metaphor on and on the evaluative load. Also, the framework captures referential uses (according to which a slur is used to mean "friend", "buddy"", "man") by foregrounding AGENTIVE – the meaning dimension

---

[5] As with the meaning dimensions of "book", more than one can be foregrounded when slurs are used. In fact, it might be argued that both EVALUATIVE and ORIGIN need to be selected when a slur is used derogatorily, because ORIGIN gives information about the group that the speaker harbours a negative evaluation towards, thus identifying which group it is that the slur is applied to. I leave this issue, as well as many others that arise in this connection, aside.

containing information about how those the slur applies to came to be and the type of beings they are (that is, humans), while the other dimensions (especially EVALUATIVE, since this is a neutral use) are backgrounded. This allows one to predicate *merely* being (a fellow) human about the members of the target group – that is, allows them to use the slur referentially.[6]

4. Croom's view

In this section and the next I present two views that are similar to the rich-lexicon application presented above, known in literature as "stereotype-theories": Adam Croom's bundle of properties view and Eli Neufeld's essentialist theory of slurs.[7] Neither of the two authors take their view to fall under the "rich-lexicon" label (at least not explicitly), and neither of them takes slurs to be polysemous (again, at least not explicitly). However, as I hope it will become clear in what follows, the similarities between these and the view presented in the previous section are significant and thus mandates grouping them under the same label.

In a series of papers, Croom (2011, 2013, 2014, 2015, 2018) has put forward a view of slurs according to which they encode rich conceptual structures comprised of *properties*, both positive and negative, that the stereotypical members of the target groups are taken to possess. These properties correspond to the stereotype that members of the target group are thought to embody in a certain community. Thus, using multiple sources about the use of the slur "spic", Croom proposes the following list of properties associated with it that together form the lexical entry of the slur:

> "Spic"
> A1. x is Mexican American.
> A2. x is a foreign worker or exchange student with a thick non-native accent.
> A3. x is poor, from a low-income family, or engages in subservient work.
> A4. x is commonly the recipient of poor treatment.
> A5. x is very passionate, sexually suave, and family oriented.

---

[6] As I made clear in footnote 2, I assume that referential and identificatory uses are two different types of uses of slurs.

[7] Two views that could be said to be similar to the ones I present but which I won't engage with here are Ritchie's (2017) and Miščević's (2016). Ritchie explicitly claims that slurs (or at least the one she focuses on, "bitch") are polysemous, but her account is not as fine-grained as the ones I present, and she doesn't connect it to theories from lexical semantics. Miščević puts forward a view according to which the meaning of slurs is quite rich, but he lumps together both semantic and pragmatics features, thus offering a comprehensive profile of slurs, not one concerned exclusively with their lexical/semantic meaning. Additionally, Miščević's view is multipropositionalist, which avoids questions about which meaning dimensions are selected in a given context by replacing them with questions about which propositions are expressed in a given context.

A6. x is independent, hardworking, and tenacious.

A7. x is loud and out of control.

A8. x is manipulative or involved in illegal activities.

A9. x is good at singing and dancing.

A10. x is gullible, unintelligent, or naive. (Croom, 214: 162-163)[8]

However, persuaded by work about how children classify items as belonging to the same category (e.g., Rosch & Mervis (1975)), Croom holds a Wittgensteinean *family resemblance* conception of category membership: there is no essential property that all members of the target group need to share for the slur to apply, only that different members share some of the properties associated with them. So, as in the case of a particular being not having to possess all the properties stereotypically associated with dogs (e.g., having four legs, having a tail, barking, chasing cats, etc.) to qualify as being applied the label "dog", a certain individual doesn't have to possess all the properties on the list above to qualify as being the target of the slur "spic". That allows the application of the slur to individuals that are not Mexican American (a result Croom takes to make his view better than many rival ones) – and, in general, the application of the slur to individuals that don't possess any of the properties listed. In other words, none of the properties on the list are necessary properties that an individual has to possess in order to be the target of the slur; it suffices that the members of the target group share some of them in certain contexts.

How does one communicate with slurs on this picture? According to Croom, the properties that comprise a slur's conceptual structure are ranked: those ranked higher (e.g., A1) are more salient than those ranked lower (e.g., A6) in indicating category membership. This ranking, however, ca be overturned, and a speaker can select among the properties in question those which are most suitable (strategically apt) for a given communicative situation, while overlooking others. So, for example, when a slur is used derogatorily, the negative properties stereotypical members of the target group are taken to have are selected; in Croom's terms, when one uses "spic" derogatorily, some among the properties A1, A2, A3, A7, A8, A10 are selected, the speaker thus employing "spic" the as "the category that most efficiently and economically predicates the intended properties of their target and most forcefully expresses a negative attitude towards them, at least to an extent that is better than other categories available

---

[8] Similar lists are given for slurs targeting Italian Americans ("guido") in Croom (2015) and Chinese Americans ("chink") in Croom (2018).

to the speaker in their lexical inventory" (Croom 2011: 357). In contrast, when a slur is used in an appropriated manner, the positive properties stereotypical members of the target group are taken to have are selected – for example, some among the properties A1, A4, A5 A6, A9 on Croom's list. Croom is silent on the exact nature of the mechanism that selects the properties in question, but one can fill this gap by appealing to any of the mechanisms mentioned above in section 2 – or, alternatively, use the placeholders foregrounding and backgrounding, as I did.

There is thus a lot in common that Croom's view has with the one presented in the previous section: both appeal to various elements of meaning (properties, for Croom; meaning dimensions for Zeman (2022)), that stand in certain relations and together comprise the literal, lexical meaning of a slur, and both appeal to selection mechanisms to explain how a particular sense of the slur is conveyed in a certain context. However, I think that Croom's view suffers from a few weaknesses that the one I favor avoids, and I expand on these in what follows.

One issue that could be raised is that Croom doesn't consider referential or identificatory uses of slurs, and it is not entirely clear how his view handles them.[9] Regarding the latter, one can extrapolate from the above how they would be treated: presumably, when a slur is used in an identificatory manner, none of the positive or the negative properties from the list are selected; at minimum, what would remain is the property that the member of the target group the slur is predicated about is a member of the relevant group (so, for the slur "spic", A1 on the above list: x is Mexican American). Accounting for the former type of use, however, is more problematic for Croom: since the lexical entry of a slur is made out of properties none of which predicates being human, or being a "buddy", about members of the target group (they do predicate being Mexican American, as we saw, as well as a host of other properties), it is not clear how one could use a slur to communicate precisely that (that is, use the slur referentially).[10] Note that this is not an issue for the view I propose: as already mentioned, identificatory uses are accounted for by foregrounding ORIGIN – the meaning dimension containing information about a group's origin, history, and social standing, while referential uses are accounted for by foregrounding AGENTIVE, the meaning dimension containing information about how those the slur applies to came to be and the type of beings they are (that

---

[9] Once again (see footnotes 2 and 6), I stress that I consider referential and identificatory uses to be two different types of uses of slurs. If one is not on board with this claim, then Croom can treat referential uses in the same vein as identificatory ones, and thus the problem described in what follows doesn't arise.

[10] It could be thought that by predicating being Mexican American about members of the target group *entails* predicating being human beings about them, but what kind of phenomenon is this entailment? If it is a pragmatic one, then the property of being human doesn't enter the lexical meaning of the slur; if it is a semantic phenomenon, then the property should be part of the list Croom offers. A semantic explanation of the various uses of slurs offers a unitary and thus more economic treatment and should thus be preferred over a mixed (some uses receive a semantic treatment, others a pragmatic treatment) one.

is, humans). This allows one to predicate being (a fellow) human about the members of the target group, thus amounting to an advantage of my view over Croom's.

Another issue concerns Croom's reliance on stereotypes. This is a criticism that has been previously levelled against views like Croom's (see, for example, Jeshion's (2013b) forceful objections against stereotype-based views). Generally, the worry with stereotype-based views is that one can use a slur – especially with the aim to derogate, in what Jeshion calls "weapon-uses" – without knowing much (or even anything) about the properties members of the target group (are taken to) instantiate. One can thus use a slur as a weapon simply based on a negative evaluation of (or attitude towards) the relevant group. The issue is not that a certain property get selected in Croom's view as relevant for communicative purposes in a certain context but not in another (this is what allows him to account for appropriated uses, for example), but that one can in fact derogate by using a slur even if that certain property doesn't get selected in *any* context (and, in fact, need not be selected) – this holding for each property on the list, of course. Conversely, reliance on stereotypes is not necessary for appropriated uses of slurs either (and presumably neither for referential uses, nor for identificatory ones): appropriating a slur can, but need not, be conditioned on reliance on a stereotype about the target group. The primary aim of appropriation many times is (or can be) to revert the negative evaluation of the members of the target group, and not to change the stereotype – even if, and possibly quite often, as a result of appropriation, the stereotype changes as well.[11] Additionally, as Jeshion points out, "[e]ven when a stereotype explains why a bigot takes a derogatory attitude toward a group, the stereotype need not be semantically encoded in the slur. More generally, it is not clear what explanatory advantage is secured here by positing semantically expressed stereotypes. Everything that needs to be explained about the speaker who is ignorant of (…) stereotypes [about a certain social group] can be done by a theory without semantically encoded stereotypes." (Jeshion 2013b, 320).[12] Importantly, on the view I advocate, since stereotypes are not part of lexical entry of a slur (and thus need not be grasped by a user of a slur for it to have its intended effects), this problem doesn't arise. Some of the properties on Croom's list might make it in the meaning dimensions that the view postulates – e.g., the property of belonging to a certain social group), but they are not part of a stereotype associated

---

[11] Croom claims that "the properties that the speaker endorses the expression of a negative attitude towards are properties that have been associated with members of a particular racial group, and as a result, the speaker does not directly express a negative attitude towards the agent him or herself" and "an agent might indirectly express a negative attitude towards a target by expressing a negative attitude towards some set of properties that that target possesses" (Croom 2011: 353). The considerations above militate against both these claims.
[12] See, though, Neufeld (2019, section 1.3) for a reply to this kind of objection.

with the target group (in the sense Croom uses the term "stereotype", as encapsulating various beliefs commonly held in a community).

Another potential problem for Croom's view is the length of the list of properties making up the lexical entry of slur proposed. As he himself acknowledges, the list can in some cases be quite lengthy, and perhaps even open-ended. This is sensible, given the fact that the list relies on information provided by stereotypes, which are usually complex and multifaceted. However, one might worry not only about whether knowing such a list is necessary for grasping the full meaning of a slur, but also whether doing so imposes a significant cognitive load vis-à-vis learning and processing information related to slurs. As already stressed above, it seems intuitive that one can use a slur (especially to derogate) without having knowledge of any of the properties associated with the target group. Importantly, in the framework I propose, this issue doesn't arise: the meaning dimensions postulated are few and comprise only essential information. While "rich", the structure of the lexicon advocated is quite minimal and the information relatively scarce – at any rate, less complex than the list of properties advanced by Croom.

Another problem that Croom's view seems to have concerns appropriation. Croom has included both negative and positive, as well as neutral, properties on the list he offers for the lexical entries of various slurs, and when one uses a slur in an appropriated manner, the positive instead of the negative properties from that list are selected. This works, of course, only if on the list of properties there are some positive traits to be selected. But one can imagine the stereotype associated with a slur – and thus the list of properties in its lexical entry – as being comprised only of negative properties (for example, it is hard to imagine any genuinely positive property associated with "nazi", should that term be considered a slur). But, if so, appropriation won't be possible in Croom's view. Yet it seems that appropriation should be possible – perhaps even more stringently so, in the case of subordinated and oppressed groups – when the stereotype associated with a slur is comprised of negative properties only (this might be the case even with the term "nazi" – say, in a dystopian future where the politics of the extreme right becomes mainstream and its adherents aim to appropriate the slur). What seems to be missing in Croom's view is the *evaluative* element, as separated from the descriptive properties that make up the lexical entries of slurs, and which is changed in appropriation. Croom claims that "[s]lurs have mixed content in the sense that the use of (…) slurs (…) can be analytically decomposed into both expressive and descriptive aspects. [B]y choosing to use the slur (…) instead of a neutrally descriptive term (…), the speaker intends to express (i) their endorsement of a (usually negative) attitude (ii) towards the descriptive properties possessed by the target of

their utterance" (Croom 2011: 353). However, the negative attitude mentioned in (i) is *not* part of the list Croom submits as the lexical meaning of a slur and is thus not part of its semantics.[13] In contrast to Croom's view, the evaluative element is clearly present in the meaning of slurs (via the EVALUATIVE meaning dimension) in the view I sketched in section 3.

Finally, an additional issue of possible concern is related to the selection of properties in a certain context, which in turn determines the way in which a slur is used. We have seen that Croom claims that the selection in question serves communicative purposes. He also submits that it is based only on closeness to the members of the target group. For example, a speaker uses a slur in an appropriated manner in a certain context only if the speaker is in a close relationship with the members of the target group. This is not an unreasonable idea: many authors hold that the permissibility or felicity of using a slur in such a way depends precisely on the tight relationship holding between the speaker and the group in question.[14] Yet, there are cases in which the use of a slur is restricted either by the linguistic material of the sentences in which it appears or by the type of context. For example, it would be entirely odd for a speaker to use a slur in an appropriated manner in a sentence stating, say, the results of a census – even if the speaker is very close to the members of the target group.[15] While this could be said to be an objection that doesn't touch the core of Croom's rich-lexicon proposal, the fact that there is an issue in this part of his view as well undermines the viability of the full package proposed.

5. Neufeld's view

Neufeld (2019, 2022) proposes an "essentialist" theory of slurs, according to which they are "failed kind terms" – that is, terms with null extension. She claims that slurs are to be treated semantically on a par with natural kind terms like "water", "gold", or "tiger", which according to her "encode an essence of a kind, *k*, that is explanatorily connected to a set of stereotypical features associated with *k*" (Neufeld 2019: 1). However, "[s]lur terms are distinctive because they designate an essence that is explanatorily connected to a set of *negative* stereotypical features of a social group" (Neufeld 2019: 2, her emphasis; importantly, "designate" is not taken by Neufeld to be a success term). According to Neufeld, slur concepts are made out of

---

[13] Croom could claim that the evaluative element constitutes a *separate* layer of meaning, as many proponents of semantic theories do. However, this is not how the view is spelled out in his works (despite the quote above).

[14] More recently, Cepollaro (2020) and Cepollaro and Lopez de Sa (2022) talk about speakers being given the *authority* to use a slur in an appropriated way by members of the target group (echoing an idea already present in Anderson and Lepore (2013a)).

[15] The situation might be different if the speaker is part of the target group. Here, however, we are looking at cases in which the speaker, although being in a close relationship with the members of the target group, is not himself or herself a member of it.

three core elements: the essence (the "underlying natures that make them the thing that they are" (Medin 1989: 1476)), the negatively evaluated stereotypical properties of the target group, and the law-like connection between the first two elements that assures that the stereotypical properties members of the target group are taken to have are explained by the essence. Here is a useful schema that Neufeld provides (2019: 5, Figure 1):
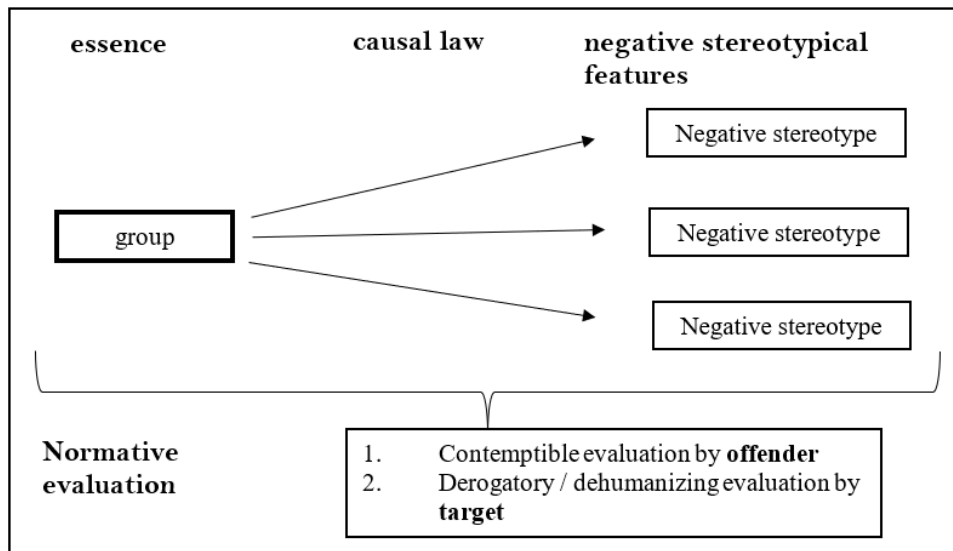

Figure 1. Model of lexical entry of a slur.

For Neufeld, then, "slur concepts encode minitheories which represent an essence-like element that is causally connected to a set of negatively-valenced stereotypical features associated with a social group" (Neufeld 2019: 3). The essences in question are such that, for example, the n-word designates a "blackness essence", "faggot" designates a "gay essence", and so on.

This is not the end of the story, however. The novelty of Neufeld's view in this context is that, in contrast to natural kind terms, slur terms have *null extension* – and that is because the essences at stake ("blackness", "gayness", etc.) don't exist. Since there are no such essences, even though they figure in the concepts people have when using slurs, the slurs themselves fail to refer.[16]

Neufeld's view has obvious commonalities with Croom's in that both postulate rich lexical entries for slurs containing properties and that both rely on stereotypes. However, in

---

[16] This is, of course, not the only view on the market that postulates null extensions for slurs: Hom (2008) and Hom and May (2013) are precursors. However, the reason slurs have null extension in their view (roughly, that there is no normative property that a person should be treated in a negative way due to their group membership) is different from Neufeld's reason (there is no essence of the target group that grounds the negative evaluation of its members).

Neufeld's proposal, the elements of these entries are related in a precise way (the essences cause the negatively evaluated stereotypical properties of the target group) instead of being more like a list. Most importantly, as we have just seen, for Neufeld slurs have null extensions (due to the inexistence of the target groups' essences), while for Croom they don't.

I have a lot of sympathy for Neufeld's view, and I agree that it explains very well derogation, and in a way that is consistent with a large portion of evidence from cognitive psychology (see also Neufeld (2022)). However, despite her claim that "[g]ood theories should predict and account for a *wide range* of data" (Neufeld 2019: 23, her emphasis) – and indeed, the rage of data she considers is impressive – Neufeld doesn't consider neither appropriated nor neutral uses of slurs. Accounting for those in her framework doesn't strike me as entirely straightforward, though. In fact, I think there is a serious problem with extending the view to such uses, which is the following. According to the proposal, slurs encode an essence that causes the negative properties ascribed to the targets; since such essences don't exist, slurs have null extension. However, slurs don't seem to have null extensions when used in an appropriated, referential or identificatory manner. When used in that way, they manage to convey positive or neutral evaluations *about* the target group, which is undoubtedly taken to exist. This raises at least a puzzle.

I see two possible replies on Neufeld's behalf. One is to claim that, in contrast to derogatory uses, when speakers use slurs by entertaining the postulated essence, this doesn't happen with the other, non-derogatory uses. In other words, the postulated essence simply disappears when a slur is not used derogatorily. This reply is blatantly ad-hoc. Alternatively, Neufeld could respond that the essence still exists, but it is now taken to be the cause of *positive* properties that the speaker, by using the slur in an appropriated manner, associates with the target group. This is a better reply, but ultimately unsatisfactory. First, what is required is a story about how the essence has gotten to be taken to be the cause of positive properties, after it has been taken, widely and for a long time, to be the cause of negative ones. It is true that, in appropriating a slur, there is a conscious effort on behalf of speakers appropriating it to re-semanticize the slur and reverse the evaluation, so that the entire process at least starts with someone's intention to use a word differently than it has been used before. But, while this story relating speakers' intention with a change in the meaning of a slur is a relatively simple one (although by no means easy to account for), a detour through essences, as being the ones responsible for the change in the slur' meaning, seems to complicate matters. In other words, appropriation aims at changing the evaluation of a target group, not its essence. A similar worry arises in connection to neutral uses of slurs: keeping the postulated essence and taking it to

explain *neutral* properties (such as being a human or a "buddy", when a slur is used referentially or as belonging to a certain social group, when a slur is used in an identificatory manner) seems like an unnecessary complication. Secondly, and still in connection to neutral uses, having the essence explain neutral properties when slurs are used neutrally would, in fact, not explain much: the essence is basically explaining that members of the target group are humans or that they belong to a certain social group. Besides, by going this route, Neufeld won't be able to differentiate between slurs (used neutrally) and their neutral counterparts – something she takes to be a good feature of her view. Given that the other option, of having the essence disappear without also providing an explanation of why this happens in appropriated and neutral uses of slurs, is ad-hoc, there seems little room for maneuverer for Neufeld to explain such uses in her framework – at least as it currently stands. In contrast, the theory I support has no problems with accounting with appropriate and neutral uses of slurs; indeed, it has been devised precisely with the aim of accounting for them.

There might be another, perhaps more direct objection to Neufeld's postulation of essences in the lexical entry of slurs. She cites a lot of psychological evidence for essences, according to which essentialization is a core characteristic of humans' conceptual ability. This is a very commendable trait of the view and offers solid theoretical support. However, one could wonder whether her view allows for competent users of a slur who are not committed to, or who even explicitly reject, any essences. Thus, imagine a social constructivist who thinks that race is not something people essentially possess, but determined by the milieu in which one grows up. In fact, they might even hold that the traits associated with, say, a given race are not immutable, but that, as a matter of statistics, the traits that most members of a certain racial group possess are negative. Now, as it happens, our social constructivist is also a racist who thinks that the traits associated with the race in question makes members of that group inferior to others and despises them. When they use a slur, they seem to be employing it without appeal to any essence. Such racist social constructivists seem conceivable. Can Neufeld make room for them?

One reply on behalf of Neufeld is to say that the racist social constructivist envisaged doesn't have full competence with the slur. But I'm not sure what the basis of such a claim would be. Neufeld cannot claim that the racist social constructivist is not competent with the slur because she doesn't admit of essences; that would be question begging. What other senses are there in which the envisaged character be not fully competent in using the slur? After all, they apply it to the right target group, and they hold the "right" (i.e., negative) attitudes towards its members. Wherein lies the mistake, then?

A second reply is that I have misconstrued the sense of "essence" she employs: in her view, essences, although fictional in the case of slurs and leading to them having null extensions, are psychologically real in the sense that they guide people's categorization of entities in their surroundings. Note, however, that the example I'm putting forward doesn't rely on the main character being a sceptic or an eliminativist about essences *tout court*: our racist social constructivist can easily be imagined to accept essences for other types of objects (such as water, gold or tigers), but not for objects pertaining to the social realm, among which race, gender and the like seem to belong. Of course, one could think the example is impossible because no one – including our racist social constructivist – can lead their cognitive life by categorizing objects from the social realm without appeal to essences, despite their assurance that they doesn't believe in essences. In this case, one can reply by making the racists social constructivist also an expert in cognitive psychology that simply has a different view on the matter. I'm not sure whether this example remains convincing, and whether making it so won't ultimately require arguing for an internalist account of semantic/conceptual competence, but I think it makes it clear that my objection is not based on a misconstrual of Neufeld's use of "essence". Needless to say, this issue doesn't appear for my preferred view, which eschews appeal to essences altogether.

6. Conclusion: a quick comparison

Where does this leave us? I have presented in some detail three views on slurs that I claimed all belong to the family of rich-lexicon theories (despite not all of them identifying as such) – views that postulate rich conceptual structures in the lexical entries of slurs, among which the sense of the slur is selected via a selection mechanism. The fact that there is more than one developed theory that can rightly be claimed to belong to this family is in itself an encouraging fact and stands to witness the fruitfulness of the framework. So, one way to interpret my considerations in this paper is as showcasing the richness of the rich-lexicon framework, and as drawing an ecumenical conclusion. While it is true that many details need to be sorted out and some objections answered in relation to each of the three views, the application of the rich-lexicon framework to slurs seems to have the potential to become a promising research program.

But my aim was also critical. What I attempted to claim is that the theory I prefer (developed in Zeman (2022) and presented in some detail in section 3) fares better than the other two with respect to several objections. For example, it fares better than Croom's in that it has a ready explanation not only of appropriated and identificatory uses of slurs, but of

referential ones as well (under the assumption that they are different), it doesn't rely on stereotypes and thus doesn't have to answer the various objections levelled against stereotype-based views, it is less complex and thus doesn't raise learning and information processing issues, it postulates a genuine evaluative element in the semantics and it doesn't tie the selection of senses to the closeness relation between speakers and members of the target group. In relation to Neufeld's view, I have complained that her view is incomplete because non-derogatory uses of slurs are not considered and that it is not clear how, as the theory currently stands, she could account for them. I also presented some misgivings about postulating essences in the lexical entry of slurs by means of an (admittedly far-fetched) example; importantly, none of these complaints appear for the view I prefer. Although there might be other, extant or potential, better situated rich-lexicon theories of slurs, the conclusion of this paper is that the one developed in Zeman (2022) emerges as the superior one to the other two brethren of the rich-lexicon family considered.

Acknowledgments

References

Anderson, L. (2018). Calling, Addressing, and Appropriation. In Sosa (2018), 6-26.

Anderson, L. & Lepore, E. (2013a). Slurring Words. *Noûs* 47(1), 25-48

Anderson, L. & Lepore, E. (2013b). What Did You Call Me? *Analytic Philosophy* 54(3), 350-363.

Asher, N. (2011). *Lexical Meaning in Context: A Web of Words*. Cambridge University Press.

Bianchi, C. (2018). Perspectives and Slurs. In G. Preyer (Ed.), *Beyond Semantics and Pragmatics*, Oxford University Press, 187-198.

Bianchi, C. (2014). Slurs and appropriation: An echoic account. *Journal of Pragmatics* 66, 35-44.

Bolinger, R. J. (2017). The Pragmatics of Slurs. *Noûs* 51(3), 439-462.

Brontsema, R. (2004). A Queer Revolution. *Colorado Research in Linguistics* 17(1), 1-17.

Burnett, H. (2020). A Persona-Based Semantics for Slurs. *Grazer Philosophische Studien* 97(1), 31-62.

Camp, E. (2018). A Dual Act Analysis of Slurs. In Sosa (2018), 29-59.

Camp, E. (2013). Slurring Perspectives. *Analytic Philosophy* 54(3), 330-349.

Cappelen, H. & Hawthorne, J. (2009). *Relativism and Monadic Truth*. Oxford University Press.

Carston, R. (2021). Polysemy: Pragmatics and sense conventions. *Mind and Language* 36(1), 108-133.

Castroviejo, E., Fraser K, & Vicente, A. (2021). More on Pejorative Language. *Synthese* 198, 9139-9164.

Cepollaro, B. (2020). S*lurs and Thick Terms: When Language Encodes Values*. Lexington Books.

Cepollaro, B. (2017a). When evaluation changes. *Journal of Pragmatics* 117, 29-40.

Cepollaro, B. (2017b). Let's Not Worry About the Reclamation Worry. *Croatian Journal of Philosophy* 17(2), 181-194.

Cepollaro, B. (2015). In Defense of a Presuppositional Account of Slurs. *Language Sciences* 52, 36-45.

Cepollaro, B. & Lopez de Sa, D. (2022). Who Reclaims Slurs? *Pacific Philosophical Quarterly*.

Cepollaro, B. & Stojanovic, I. (2016). Hybrid Evaluatives. *Grazer Philosophische Studien* 93, 458-488.

Cervone, C., Augoustinos, M. & Maass, A. (2021). The Language of Derogation and Hate: Functions, Consequences, and Reappropriation. *Journal of Language and Social Psychology* 40(1), 80-101.

Croom, A. (2018). Asian slurs and stereotypes in the USA. A context-sensitive account of derogation and appropriation. *Pragmatics and Society* 9, 495-517.

Croom, A. (2015). Slurs and stereotypes for Italian Americans: A context-sensitive account of derogation and appropriation. *Journal of Pragmatics* 81, 36-51.

Croom, A. (2104). Spanish slurs and stereotypes for Mexican-Americans in the USA: A context-sensitive account of derogation and appropriation. Sociocultural Pragmatics 2(2), 145-179.

Croom, A. M. (2013). How to do things with slurs. *Language and Communication* 33, 177-204.

Croom, A. (2011). Slurs. *Language Sciences* 33(3), 343–358.

Del Pinal, G. (2018). Meaning, Modulation, and Context. *Linguistics and Philosophy* 41(2), 165-207.

Diaz Legaspe, J., Liu, C. & Stainton, R. (2020). Slurs and register. *Mind & Language* 35(2), 156-182.

DiFranco, R. (2014). Pejorative Language. *Internet Encyclopedia of Philosophy*.

Falkum, I. L. & Vicente, A. (2020). Polysemy. *Oxford Bibliographies Online*. Oxford University Press.

Falkum, I. L. & Vicente, A. (2015). Polysemy: Current Perspectives and Approaches, *Lingua* 157, 1-16.

Frisson, S. (2009). Semantic Underspecification in Language Processing. *Language and Linguistics Compass* 3(1), 111–127.

Gutzmann, D. (2019). *The Grammar of Expressivity*. Oxford University Press.

Gutzmann, D. (2015). *Use-conditional Meaning. Studies in Multidimensional Semantics*. Oxford University Press.

Hess, L. (2020). Practices of Slur-Use. *Grazer Philosophische Studien* 97(1), 86-105.

Hogeweg, L. & Vicente, A. (2020). On the Nature of the Lexicon. *Journal of Linguistics* 56(4), 865-891.

Hom, C. (2012). A puzzle about pejoratives. *Philosophical Studies* 159, 383-405.

Hom, C. (2010). Pejoratives. *Philosophy Compass* 5/2, 164-185.

Hom, C. (2008). The Semantics of Racial Epithets. *The Journal of Philosophy* 105(8), 416-440.

Hom, C. & May, R. (2018). Pejoratives as Fiction. In Sosa (2018), 108-131.

Hom, C. & May, R. (2013). Moral and Semantic Innocence. *Analytic Philosophy* 54(3), 293-313.

Hornsby, J. (2001). Meaning and uselessness. *Midwest Studies in Philosophy* 25, 128-141.

Jackendoff, P. (1990). *Semantic Structures*. MIT Press.

Jeshion, R. (2020). Pride and Prejudiced: On the Reclamation of Slurs. *Grazer Philosophische Studien* 97(1), 106-137.

Jeshion, R. (2018). Slurs, Dehumanization, and the Expression of Contempt. In Sosa (2018), 77-107.

Jeshion, R. (2016). Slur Creation, Bigotry Formation: The Power of Expressivism. *Phenomenology and Mind* 11, 130-139.

Jeshion, R. (2013a). Expressivism and the Offensiveness of Slurs. *Philosophical Perspectives* 27, 231-259.

Jeshion, R. (2013b). Slurs and Stereotypes. *Analytic Philosophy* 54(3), 314-329.

Jusińska, Z. (2021). Slur Reclamation – Polysemy, Echo, or Both? *Organon F* 28(3), 689-707.

Kennedy, R. (2003). *Nigger: The Strange Career of a Troublesome Word*. Penguin Random House.

Kirk-Giannini, C. D. (2019). Slurs are Directives. *Philosophers' Imprint* 19(48).

Langton, (2012). Beyond Belief: Pragmatics in Hate Speech and Pornography. In I. Maitra & M. K. McGowan (Eds.), *Speech and Harm: Controversies Over Free Speech*. Oxford UP, 72-93.

Liu, C. (2021). The Derogatory Force and the Offensiveness of Slurs. *Organon F* 28(3), 626-649.

Marques, T. & García-Carpintero, M. (2020). Really Expressive Presuppositions and How to Block Them. *Grazer Philosophische Studien* 97(1), 138-158.

McCready, E. (2010). Varieties of conventional implicature. *Semantics and Pragmatics* 3(8), 1-57.

Medin, D. (1989). Concepts and Conceptual Structure. *The American Psychologist* 44(12), 1469-1481.

Miščević, N. (2016). The Fiery Tongue: The Semantics and Pragmatics of Pejoratives. In N. Miščević & J. Perhat (Eds.), *The Word Which Bears a Sword. Semantics, Pragmatics and Ethics of Pejoratives*, Kruzak, 11-156.

Naylor, G. (1986). The Meanings of a Word. *The New York Times*.

Neufeld, E. (2022). Psychological essentialism and the structure of concepts. *Philosophy Compass* 17(5), e12823.

Neufeld, E. (2019). An Essentialist Theory of the Meaning of Slurs. *Philosophers' Imprint* 19(35), 1-29.

Nunberg, G. (2018). The Social Life of Slurs. In D. Fogal, D. W. Harris & M. Moss (Eds.), *New Work on Speech Acts*, Oxford University Press, 237-295.

Orlando, E. & Saab, A. (2020). Slurs, Stereotypes, and Insults. *Acta Analytica* 35, 599-621.

Ortega Andrés, M. & Vicente, A. (2019) Polysemy and Co-predication. *Glossa* 4(1).

Popa-Wyatt, M. (2020). Reclamation: Taking Back Control of Words. *Grazer Philosophische Studien* 97(1), 159-176.

Popa-Wyatt, M. & Wyatt, J. (2018). Slurs, Roles and Power. *Philosophical Studies* 175(11), 2879-2906.

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

Potts, C. (2007). The Expressive Dimension. *Theoretical Linguistics* 33(2), 165-198.

Predelli, S. (2013). *Meaning without Truth*. Oxford University Press.

Pustejovsky, J. (1995). *The Generative Lexicon*. Cambridge, MA: MIT Press.

Richard, M. (2008). *When Truth Gives Out*. Oxford: Oxford University Press.

Ritchie, K. (2017). Social Identity, Indexicality, and the Appropriation of Slurs. *Croatian Journal of Philosophy* 17(2), 155-180.

Romanian Government Report (2009*). Barometrul relaţiilor interetnice. Mai-iunie 2009. Etnici romi – raport parţial*. (*Interethnic Relations, May-June 2009. Roma ethnics, partial report*).

Rosch, E. & Mervis, C. (1975). Family resemblances: studies in the internal structure of categories. *Cognitive Psychology* 7, 573-605.

Saka, P. (2007). *How to Think about Meaning*. Springer.

Schlenker, P. (2007). Expressive presuppositions. *Theoretical Linguistics* 33, 237-245.

Schumacher, P. (2013). When Combinatorial Processing Results in Reconceptualization: Towards a New Approach of Compositionality. *Frontiers in Psychology* 4, 677.

Scott, M. & Stevens, G. (2019). An Indexical Theory of Racial Pejoratives. *Analytic Philosophy* 60(4), 385-404.

Smitherman, G. (2006). *Word from the mother. Language and African Americans*. Routledge.

Sosa, D. (Editor). *Bad Words. Philosophical Perspectives on Slurs*. Oxford University Press.

Stojnić, U. & Lepore, E. (2022). Inescapable articulations: Vessels of lexical effects. *Noûs* 56, 742-760.

Sullivan, A. (2022). Semantic Dimensions of Slurs. *Philosophia* 50, 1479-1493.

Tirrell, L. (1999). Derogatory Terms: Racism, Sexism and the Inferential Role Theory of Meaning. In K. Oliver & C. Hendricks (Eds.), *Language and Liberation*. SUNY Press, 41-79.

Vallée, R. (2014). Slurring and common knowledge of ordinary language. *Journal of Pragmatics* 61, 62-87.

Vicente, A. (2021) Approaches to co-predication. *Journal of Pragmatics* 182, 348-357.

Vicente, A. (2018). Polysemy and Word Meaning. *Philosophical Studies* 175(4), 947-968.

Vicente, A. (2010). Clusters: On the Structure of Lexical Concepts. *Dialectica* 64(1), 79-106.

Vicente, A. & Falkum, I. L. (2017) Polysemy. In M. Aronoff (Ed.), *Oxford Research Encyclopedia of Linguistics*. Oxford University Press.

Whiting, D. (2013). It's Not What You Said, It's the Way You Said It. *Analytic Philosophy* 54(3), 364-377.

Whiting, D. (2007). Inferentialism, Representationalism and Derogatory Words. *International Journal of Philosophical Studies* 15(2), 191-205.

Williamson, T. (2009). Reference, Inference, and the Semantics of Pejoratives. In J. Almog & P. Leonardi (Eds.), *The Philosophy of David Kaplan*. Oxford University Press, 137-158.

Zeman, D. (2022). A Rich-Lexicon Theory of Slurs and Their Uses. *Inquiry* 65 (7), 942-966.