



Hogeweg, Lotte. 2023. Dutch slurs and the descriptive-expressive distinction. *Glossa: a journal of general linguistics* 8(1). pp. 1–18. DOI: <https://doi.org/10.16995/glossa.9246>



Open Library of Humanities

Dutch slurs and the descriptive-expressive distinction

Lotte Hogeweg, Radboud University, Centre for Language Studies, NL, lotte.hogeweg@ru.nl

There is a lot of debate about the meaning of slurs, which are words that refer to a certain group of people in a derogative manner. Many scholars (e.g. McCready 2010; Croom 2011; Gutzmann 2015) consider slurs, unlike their 'neutral' counterparts, to have both descriptive and expressive content. In this squib, I discuss initial data on the use of Dutch slurs and related terms on Twitter. The Dutch data show that 'neutral' labels sometimes seem to behave like slurs to some degree and that some slurs can be used as swear words without (seemingly) conveying any descriptive property. Based on these observations, I suggest that slurs may be part of a continuum of labels, ranging in the degree to which they have descriptive versus expressive content, entailing a gradual distinction between conversational and conventional implicatures in the lexical domain. A gradual view on conventionality in the lexical domain would moreover pave the way for a less fundamental distinction between expressive and social meaning. It is concluded that the debate on slurs would benefit from a broader and more empirical approach to pejorative terms for (groups of) people, as their actual use and interpretation can show a mixture of (degrees of) descriptive, expressive and social meaning.

Glossa: a journal of general linguistics is a peer-reviewed open access journal published by the Open Library of Humanities. © 2023 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <http://creativecommons.org/licenses/by/4.0/>.

OPEN ACCESS



1 Slurs and the distinction between expressive and descriptive content

Traditionally, semanticists, at least formal semanticists, focused on language as a means to describe the world as in (1). Kaplan (1999) and Potts (2005) were among the first semanticists who elaborated on the idea that language can also be used not to describe the world, but to express one's attitude toward an aspect of the world, as the word *idiot* does in (2).

- (1) John forgot his umbrella.
- (2) That **idiot** John forgot his umbrella.

Where Potts argued that words or other linguistic elements have either descriptive content or expressive content, scholars thereafter, such as McCready (2010) and Gutzmann (2015), have shown that there are elements that have both at the same time, that is, they have mixed content. An example is the word *cur* in (3) (Gutzmann 2019: 12, original numbering (2.1 b)).

- (3) This **cur** howled the whole night.

The word has descriptive, truth-conditional content, as it indicates that the referent is a member of the set of dogs, but at the same time, the speaker of (3) expresses a negative attitude towards dogs. Crucially, it is typically argued that there is a specific set of items that expresses both types of meaning. As Gutzmann (2019: 11) puts it, "every use of a linguistic sign may reveal something about the speaker's attitude [...]. However, there are certain expressions that have the expressive function hardwired into their lexicon".

Probably the most studied type of mixed content bearers is formed by the class of slurs. Slurs are pejorative terms for certain groups of people, like *kraut*¹ which is a derogatory term for 'German'. Slurs have mixed content because they indicate that the subject of the sentence is a member of the group the term refers to, and at the same time they express a negative attitude towards this group. There is a lot of discussion in the linguistic, (formal) semantic and philosophical literature about slurs. The discussion centers around the question what slurs convey and the status of the various types of information or content. One group of authors argues that the derogative content of slurs is part of their semantic meaning (e.g. Hornsby 2001; Hom 2008). The studies that consider slurs to express mixed content are part of this class, assuming that the slurs have an additional non-truth conditional layer of meaning that is conventionally tied to the lexical term (e.g. McCready 2010; Gutzmann 2015). Other scholars have argued that the derogative aspect is not part of the semantics of slurs, or at least not in a standard way (e.g. Camp 2013; 2018; Bolinger 2017; Nunberg 2018). In the discussion, slurs are typically contrasted

¹ I decided to include less familiar English slurs and Dutch slurs in their full form in this paper for reasons of clarity and transparency.

with their so-called neutral counterparts. Some authors (e.g. Williamson 2009; McCready 2010) argue that a slur like *n****** has the same descriptive content as its non-offensive counterpart *African American*. Others (e.g. Croom 2013; Burnett 2020) argue that slurs are not descriptively similar to their neutral counterparts but that, for example, the descriptive content consists of stereotypical negative properties associated with members of a particular group.

In this squib, I will address examples of Dutch slurs and compare them to related words. Clear exceptions aside (e.g. Carnaghi & Maass 2007; Bartlett et al. 2014; Fasoli et al. 2015; O’dea et al. 2015), the discussion on slurs is characterized by its highly theoretical nature. The purpose of the squib is to investigate the actual use of slurs and other (pejorative) terms for (groups of) people, in particular to explore whether their use reflects a categorization in descriptive terms, mixed terms and purely expressive terms. To that end, I will discuss an exploratory investigation of the use of Dutch slurs and related words on Twitter in the next section. In the subsequent discussion, I will elaborate on what the use of the terms tells us about their meaning and relate the findings to the (semantic) debate on slurs.

2 Dutch slurs – an exploratory investigation of their use on Twitter

Unlike English, Dutch does not seem to possess a large inventory of derogatory words for certain groups of people that have a clear neutral counterpart. *Mof* is a clear example, the neutral counterpart being ‘German’, and there are several slurs for homosexual people such as *flikker* or *mietje* for men or *pot* for women. Additional examples seem difficult to find, however. There are derogatory terms that are clearly related to aspects such as religion or descent, such as *geitenneuker* literally ‘goat fucker’ or *spleetoog* ‘slanted eyed’, but it is not entirely clear what the neutral counterpart of these words would be (possible candidates would be the Dutch counterparts of *Muslim* or *Arab* for *geitenneuker* and *Chinese* or *Asian* for *spleetoog*) (see Jeshion 2016 for a similar observation). In addition, there are supposedly neutral words that seem to be used as a slur sometimes such, as *Turk*, which is often used to insult people and is also used to classify people of various backgrounds other than Turkish. Moreover, the relatively clear examples of slurs for homosexuals seem to be used in other ways as well. *Mietje* for example often seems to be used not so much to indicate that a person is a homosexual but to indicate that someone has a certain negative property, something like being a coward. *Flikker* seems to often be used without any descriptive content, not to indicate that someone is a homosexual, nor to ascribe a negative property to someone, but just to express one’s discontent with someone.

The inventory of Dutch (possibly) pejorative terms for (groups of) people therefore seems to primarily consist of words whose classification would not be obvious based on a first, intuitive, consideration of how they are typically used. To substantiate these intuitions with more objective data, I collected tweets containing one of a number of possible slurs and related words to investigate how those terms are used. I selected three clearly derogative terms related to a

particular group: the aforementioned *flikker*, *mietje*, and *geitenneuker*. In addition, I selected two (supposedly) neutral words that seem to be used to insult people as well: the aforementioned *Turk* and *homo*, which is short for *homoseksueel* ‘homosexual’ and can be used in a neutral manner, while it is (often) used as a swear word as well. In addition, I selected the term *Oostenrijker* ‘Austrian’ as a comparison because this word is expected to be used in a mostly neutral way and *eikel* which is a swear word (comparable to *jerk*) that has no connection to a particular group of people (except that it is only used for men). Finally, I included the term *neger*. *Neger* is of course related to the English *n-word* but it has a different history. Until (relatively) recently its use was less loaded than the English *n-word* and by some it was considered to be a neutral descriptive term. During the last years, its use has become more controversial and it is now commonly considered a derogatory term.

In the next section, I will discuss the methods and results of the exploratory investigation. The examples that are used below to illustrate the categories and results are based on the actual tweets we found, but they are sometimes slightly adapted and also only shown as their English translation for reasons of privacy and ethics (see e.g. Townsend & Wallace 2016).

2.1 Methods and categories

With the help of Wessel Stoop of the Humanities Lab of Radboud University, at least 200 of the most recent tweets were collected for each of the terms. With the approval of Twitter, tweets were collected from an initially 30-day period before the start of the investigation (which was October 2021), extended to a longer period if the initial search did not yield enough relevant occurrences. We² manually eliminated irrelevant occurrences (e.g. occurrences where the word was not used as a noun) until we had 200 relevant uses for each of the terms.

For each term, the 200 tweets were annotated by two research assistants. Occasionally, the wider context of a tweet was taken into account (if the tweet was still available online), when the message itself did not contain enough information. At least 20 tweets per category were annotated together to sharpen the categories and at least 60 tweets were annotated by both annotators independently to calculate the inter annotator agreement. For the terms that turned out to be more difficult to classify, an extra set of occurrences was doubly annotated. All tweets were finally also considered by the author and if disagreement arose, this classification was reconsidered by the author and one of the assistants until they reached an agreement. A small number of tweets was replaced in this last round because on closer inspection they turned out to fall outside the predefined selection criteria. The newly selected tweets were considered by the author and the research assistant.

² In the discussion of the methodology, the use of *we* refers to the author and the research assistants who conducted the annotation.

The purpose of the study was to make an inventory of the ways each of the terms were used and with what frequency to explore whether a clear categorization of terms can be identified in their patterns of usage, such that, e.g. neutral terms can be distinguished from non-neutral terms. As, on the one hand, we aimed for an empirically driven annotation of the words, but on the other hand needed classification criteria in order to compare the various terms, we decided to make use of the categories proposed by Bartlett et al. (2014). Their study on the use of English slurs is one of the few empirical and quantitative studies that have been done on the topic, and as they also used Twitter as a source of data, their categories are likely to be applicable for our purposes. Since Bartlett et al. (2014) only investigated slurs and our set of terms also includes other words, however, we had to adapt and extend the categories to a certain degree. In total, we used seven categories. Two of the categories, however, occurred very infrequently and will therefore not be discussed in the remainder of the text.³ The remaining five categories are explained below.

Negative Context (NC): A term was categorized as NC if it was used in a negative context or otherwise in a negative way, as in (4). Our initial annotation included two categories that, based on Bartlett et al. (2014), distinguished between occurrences that expressed a negative stereotypical attitude and occurrences that were used in a derogative manner but without explicit stereotypical references (after Bartlett et al.'s (2014) category *Casual Use*). The distinction between the two categories seemed too fuzzy for our data in the end. A term was also classified as NC if it was used to insult a third person, as in (5). In the latter case it was often difficult to determine whether the person insulted belonged to the target group or whether the term was used as a general swear word.

- (4) A **Turk** will always be a Turk. He won't integrate. Let him run to Turkey, to Erdogan, he doesn't belong in our country.
- (5) A real **flikker**, that Eddy.

Targeted Use (TU): A term was labelled TU if the term was directly addressed to someone, as in (6). Bartlett et al. (2014) used the category Targeted Abuse, but we also categorized uses as such if the term was not necessarily used in a clearly negative manner as in (7).

- (6) You don't drink coffee? Hah, you **mietje**!
- (7) You are a **Turk**.

When a tweet consisted only of the relevant term and was directed at someone (with the @-sign), we annotated the terms as TU unless there was a reason to annotate it otherwise.

³ The two categories were *appropriated use*, i.e. when the term was used by a speaker to refer to or describe themselves, and *offline action*, which (based on Bartlett et al. 2014) meant that the writer suggested in the tweet that action should be taken in the offline world.

Non derogatory (ND): following Bartlett et al. (2014), a term was classified as ND if it was used in a descriptive or otherwise neutral way, or if a stereotype was applied but not in a hurtful or derogatory manner, as in (8).

(8) First **Oostenrijker** to win the US Open.

Mention (ME): A term was classified as ME if the term was not used but mentioned, so the term was used in a meta-linguistic manner, as in (9). This category was not applied by Bartlett et al. (2014).

(9) I do understand the word **geitenneuker** a bit better now.⁴

Negative Trait (NT): As we also included terms that did not (necessarily) refer to a particular (social) group but are rather used to ascribe a negative property to someone as in (10), we also included the category NT. The difference between NC and NT is that NT describes a trait while NC does not.

(10) A typical case of people who are tough when they are in a group but a **mietje** when they are alone.

2.2 Results

In this section, the most relevant findings and observations will be discussed.⁵ Not all numbers will be presented, and the results should be regarded with caution for reasons discussed in the next section, but I will point out a number of interesting first observations.

Let us first look at the so-called neutral words: *Oostenrijker*, *homo* and *Turk*. For reasons of readability, I only included the three most informative categories in **Figure 1**: Negative Context, Non-Derogatory and Targeted Use. Note that the occurrences were always classified in only one of the categories.

As expected, *Oostenrijker* is almost exclusively used in a neutral way (97%), as in (11). What was noticeable with respect to *Oostenrijker* was that quite frequently it was used to refer to Hitler, as in example (12). Those occurrences were also annotated as Non-Derogatory, as the term *Oostenrijker* itself does not seem to be used in a negative manner.

(11) Small correction: Gregor Mühlberger is an **Oostenrijker**.

(12) And that book by that one **Oostenrijker** with the silly hairdo and mustache.

⁴ As one of the reviewers rightfully pointed out, this instance of the category ME seems to also convey a negative attitude towards the group typically associated with the term. This is in line with more general problem of the categorization discussed in Section 3, that the categories we used seemed to be fuzzy and not mutually exclusive.

⁵ The data have been archived but, conform university guidelines, are not made publicly available for reasons of privacy and copyright.




Figure 1: The absolute frequencies of three of the categories for the ‘neutral’ terms *homo*, *Oostenrijker*, and *Turk*.

For *homo*, we already see that there are more uses in negative contexts (11,5%) and also a number of targeted uses, although there were only three of the latter category. Examples (13) to (15) illustrate the three categories.

- (13) Recognition of the mistakes made could ease the pain of **homo**’s. (Non-Derogatory)
 (14) Always something with those **homo**’s. (Negative Context)
 (15) Why do you defend pedosexuals? **Homo**! (Targeted Use)

For *Turk* we see more occurrences in a negative context (43%) and more occurrences of targeted use (8%). These are examples of each of the categories.

- (16) I just want a **Turk** with the same mind-set. (Non-Derogatory)
 (17) The prototypical **Turk** that I immediately distrust. (Negative Context)
 (18) You are still a worthless f***ing **Turk**. (Targeted Use)

In **Figure 2**, the slurs *flikker*, *mietje* and *geitenneuker*, and the swear word *eikel* are represented. For these terms, it is most informative to look at the categories Negative Context, Negative Trait and Targeted Use.




Figure 2: The absolute frequencies of three of the categories for the terms *eikel*, *flicker*, *geitenneuker*, and *mietje*.

What seems noticeable is that *flicker* and *geitenneuker* have no occurrences in the category Negative Trait. However, this is due to the fact that we did not consider this to be a possible category for these terms as it would be unclear what negative trait they would indicate (see Section 3). What is noticeable is that the supposed slurs *mietje* and *flicker* are more frequently used in a targeted manner (32,5% and 36% respectively) than the swear word *eikel* (26%).

For *flicker*, Targeted Use, as exemplified in (19) is the most frequent category.

(19) You are talking to me, aren't you, **flicker**! (Targeted Use)

The second most frequent category was Negative Context (32%). Recall that under this category we also counted cases in which a third person is insulted, called a name, as in (20).

(20) The referee is a real **flicker**. (Negative Context)

This was very frequently the case. We had the impression that most of the time, no reference was made to homosexuality both for the category Negative Context and for the category Targeted Use, but this was typically very difficult to know for sure.

For *mietje*, Negative Trait was the most frequent category (43,5%). In example (21), *Lex* refers to the Dutch king Willem-Alexander and the tweet alludes to an incident in which the king and his family flew to their summer house in Greece in the middle of a lockdown and returned to the Netherlands after nationwide criticism.

(21) Lex is a **mietje**. He should have stayed in Greece. (Negative Trait)

In such a case the term was not used to indicate that someone is a homosexual but that he has a negative quality, namely that he is cowardly. Only in a very limited amount of cases it was clear that the term did indicate that the person talked about is a homosexual, as in (22). In such cases the term was classified as NC (5%).

(22) A whiner and a **mietje**. Which is highly unusual amongst those people by the way. He hid it well though, apart from his voice that is. (Negative Context)

If the tweet was about a public figure who is known to be homosexual, we annotated the term as NC when reference was made to their homosexuality, and as NT when this was not the case. A number of tweets for example related to an incident in which an openly homosexual member of parliament pressed charges against a group of farmers who visited him at his home address as an action of protest. The word *mietje* seemed to be used as qualification of this in the user's eye cowardly act. At the same time, it is probably not coincidental that precisely this word was chosen, and sometimes explicit reference to his homosexuality was made as well. It shows that a clear distinction between categories is difficult to make, and that the categories are indeed fuzzy and not mutually exclusive. I will come back to this in the next section.

Eikel was more often used to describe someone as in (23) than to insult someone directly as in (24) (65,5% versus 26%). We categorized occurrences of *eikel* attributed to a third person as negative trait, the trait being something like rude or impertinent. I will also come back to this in Section 3.

(23) What an **eikel**, this guy. Just take down the tree, it is too big for such a small garden. (Negative Trait)

(24) You shut up, **eikel!** (Targeted Use)

What was noticeable with respect to *geitenneuker* was that it was quite often used to refer specifically to Erdoğan, the president of Turkey, as in (25). What was furthermore noticeable was that this term, more often than *mietje* and *flikker*, occurred in a context that was clearly related to the target group as in (26), and intuitively the tweets with *geitenneuker* were also the most hateful ones. Both were annotated as Negative Context (58,5%).

(25) Erdogan is a **geitenneuker**.

(26) Scary **geitenneuker**, I say it is about time that we slow down Islam.

Geitenneuker was used in a targeted manner in 21,5% of the cases, as in (27).

(27) And you're just a **geitenneuker**.

When we look at *neger*, finally, I included the categories Mention, Non-Derogatory, Negative Context and Targeted Use in **Figure 3**.




Figure 3: The absolute frequencies of four of the categories for the term *neger*.

What is noticeable is the high frequency of the category Mention (22,5%), so people were frequently talking about the word instead of using it. What is furthermore noticeable is that it is more frequently used in a negative context than in a neutral context (40,5 versus 7,5%) and that it is not so frequently used in a targeted manner (2,5%). In that sense, the word is more like *Turk* than it is like *flikker* or *mietje*. Examples (28) to (31) illustrate the four categories.

- (28) This is not a Dutchman, it is a criminal **neger**. We should deport him. (Negative Context)
- (29) She has dreads to look more like a **neger**. (Non-Derogatory)
- (30) Why do you say **neger** like you are allowed to do so? (Mention)
- (31) F***ing **neger**, shut up! (Targeted Use)

3 Discussion

Before I discuss what we can learn from these data, I would first like to address some issues or potential problems. The first is that categorization was sometimes quite difficult, as was already mentioned a number of times above. The inter-annotator agreement, computed as Cohen's Kappa for each category was between .56 and 1, with *Turk* and *homo* having the lowest score.⁶ While

⁶ Cohen's Kappa was calculated based on the annotation of the two research assistants. The calculation included the two categories Negative Stereotypical Attitude and Casual Use that were afterwards collated into the category Negative Context. The scores were as follows: *Oostenrijker*: 1, *Turk*: 0.56, *homo*: 0.59, *flikker*: 0.8, *geitenneuker*: 0.79; *mietje*: 0.74; *eikel*: 0.81, *neger*: 0.73.

this may not be too bad, it was clear that the boundaries between the categories are fuzzy and perhaps therefore not always meaningful. Furthermore, the classification was clearly informed by what we as native speakers intuitively know (or think we know) about the meaning of the words. For example, the fact that we classified an utterance as *He is a flicker* in the category Negative Context and *He is an eikel* as Negative Trait is not based on the utterances themselves but on our ideas about what the terms can mean. A term like *geitenneuker* was also never annotated as non-derogatory because our intuitions tell us that the word is never used in a neutral or playful way. We have to keep that in mind when we consider the results. A final problem of this investigation is that the data were collected from tweets from a relatively short time period. The way the terms were used and with what frequency may have been influenced by the topic that happened to be relevant on Twitter during this time period. A lot of tweets containing the term *neger* for example contained some reference to the Black Lives Matter protests.

Nonetheless, I think that even the examples already suggest that many of the terms are not easily classified. There were only two terms in our sample with a more or less clear profile: *Oostenrijker* had a clearly neutral profile, and *geitenneuker* was seemingly always used in the context of a specific group and was always used in a negative way. *Geitenneuker* would therefore be a clear example of a slur, apart from the observation mentioned above that its neutral counterpart is not obvious. We can therefore conclude that almost all words in our sample are used in a variety of ways, without suggesting a clear categorization of usage profiles. An important question, however, is how we should see the relation between the *use* of the words and their *meaning*. Does the fact that *Turk*, for example, was used more frequently in a negative than in a neutral context, mean that the word itself has a pejorative meaning? After all, almost any word can get a negative connotation in the right context. However, the frequency with which words co-occur with a certain meaning is known to be an important factor in the development of meaning. In the next section I will first briefly discuss the relation between frequency and meaning (change) and then relate this to the findings presented in Section 2.

3.1 Frequency and conventionalization

It is commonly acknowledged that words may gain expressive or social meaning and lose descriptive meaning through a process of semantic change or pragmaticalization (e.g. Foolen 1997). Gutzmann & Davis (2015) for example argue for a three-stage process in which expressions with mixed content serve as a bridge between originally descriptive expressions and their expressive counterparts. The move from stage one to stage two takes place through the conventionalization of conversational implicatures, thus becoming conventional implicatures, after which the descriptive content may be lost altogether in stage three due to semantic bleaching. It is also typically acknowledged that this process of conventionalization of pragmatic inferences is determined by frequency of use. As Gutzmann & Davis (2015: 203) put it: “Given a sufficiently high frequency, these inferences may be conventionalized and become part

of an expression's lexical content, so that they do not need to be derived on conversational grounds anymore". Schmid (2015) distinguishes between entrenchment, which takes place in individual speakers, and conventionalization, a process which takes place in communities or social groups. Schmid (2015) argues that repeated processing of identical form meaning pairings lead to routinization of symbolic association or to *context-free entrenchment*, i.e., the more a word co-occurs with a certain meaning, the less context is needed to activate this meaning. This routinization of association is a gradual process, meaning that expressive content may be entrenched and eventually conventionalized to a certain degree. During this process, "usage affects both entrenchment and conventionalization, while entrenchment and conventionalization in turn influence usage" (Schmid 2015: 7).

In the data presented in Section 2, some terms were much more frequently uttered in a negative context than others. The word *Turk*, for example, more frequently occurs in a negative than in a neutral context, while *Oostenrijker* occurs basically only in neutral contexts. What is furthermore interesting, is that it seems that the more often a term is used in a negative context, the more frequent the term is used in a targeted manner as well. *Turk*, for example, was used in a targeted manner in 8% of the cases. Utterances we attested like *Hypocritical Turk!* or *Sick Turk!* or just *Turk!*, but also, for example, *Homo!* or *F***ing neger, shut up!*, are obviously expressive. The adjectives of course add to the expressive meaning of the utterance, and the broader context (if applicable) is typically not very complimentary as well, but the question arises why the writers of such tweets would use or add the terms *Turk*, *homo* or *neger* if those terms did not by themselves convey expressive content, i.e. add to the expressive force of the utterance, as well. It seems therefore that these terms do not always depend on the context to receive an expressive interpretation. On the other end of the scale we see that *flikker* is mostly used in a targeted manner (more so than the assumed swear word *eikel*), while for example *geitenneuker* is more often used in a negative context. Following the abovementioned hypothesized relation between targeted use and expressive content, this might suggest that *flikker* is in the process of losing its descriptive meaning.

Yet, there is a clear difference between a term like *geitenneuker* or *flikker* on the one hand and a term like *Turk* or *homo* on the other hand, namely that, unlike *geitenneuker*, *Turk* and *homo* can get a (seemingly) neutral interpretation. Gutzmann (2019: 11) makes a strict distinction between terms that do and terms that do not "have the expressive function hardwired into their lexicon". A similar strict distinction between conversationally implicated and lexically expressed negative attitude is made by Camp (2013), who argues that a hearer of (32) might suspect that the speaker holds a racist attitude towards Hispanics as he finds this property of the candidate worth mentioning. She furthermore adds that the use of a noun instead of an adjective adds to this suspicion (Camp 2013: 340, original numbering (12)).

(32) They gave the job I applied for to a Hispanic.

Camp stresses, however, that such an attitude is at most implicated and can be consistently denied by the speaker. Deniability or cancellability is often taken to be a crucial characteristic that differentiates between conversational and conventional implicatures (e.g. Potts 2005). Henderson & McCready (2019: 3) argue that the deniability argument “gets at the heart of what it means for content to be conventional”. Often, deniability is motivated by means of a constructed example, as in Henderson & McCready (2019), who illustrate the non-deniability of the pejorative content of the slur *Kraut* with (33) (Henderson & McCready 2019: 3, original numbering (7)):

- (33) A: Angela Merkel is a kraut.
 B: What do you have against Germans?
 A: #I don't have anything against Germans. Why do you think I might?

The question arises whether deniability is indeed an all or nothing quality of meaning. An increasing degree of conventionalization (or routinization or entrenchment) of expressive meaning may make it harder to plausibly deny a derogative attitude towards the group referred to. A loss of descriptive content—as, for example, seems to hold for *flikker* and for *mietje* (whose descriptive content seems to narrow down to a specific property)—on the other hand, may make it easier to deny a negative attitude towards the group (previously) associated with the term. Sternau et al. (2017) experimentally show that the deniability of various types of inferences in other domains (e.g. fragment completion and conjunction enrichments) indeed seems to be a matter of degree. Whether the same holds for the deniability of pejorative content of the terms under consideration is an empirical question. If the deniability and therefore the conventionality of expressive content is indeed best viewed as a gradual property, what would be the consequences for the theoretical analysis of slurs and the related terms discussed in this paper? This will be addressed in the next subsection.

3.2 Gradual expressive meaning

If indeed the expressive meaning of a term, such as *Turk*, becomes less easily deniable and less dependent on the context, perhaps to a degree in which it is standardly activated and only cancelable or deniable in specific circumstances, how should this aspect of its meaning be characterized? One option would be to qualify it as a generalized conversational implicature. Generalized implicatures arise “in the absence of special circumstances” (Grice 1989: 37). According to Potts (2005), generalized conversational implicatures, like other conversational implicatures, differ from conventional implicatures in that they are calculable based on Grice's (1975) conversational maxims, and that, consequently, they are not non-detachable, i.e. they are not tied to a specific choice of words. In this respect it is interesting to consider the aforementioned observation by Camp (2013) that nouns may give rise to negative inferences to a stronger degree than adjectives. This suggests that there must be something inherent to their meaning that causes this. In line with Camp's (2013) suggestion, Carnaghi et al. (2008) experimentally show that people ascribe more stereotypical properties to someone when this person is described with a noun, for example as a

homosexual than if a person is described with an adjective, as *homosexual*. Hogeweg & Neuleman (2022) likewise show that some supposedly (relatively) neutral nouns are considered to be less hurtful than their corresponding slurs while they are considered to be more hurtful than their corresponding adjective noun combinations (while offered in the same, neutral context).

Recanati (2003) describes an additional type of implicature, a default implicature, that on the one hand differs from (generalized) conversational implicatures in that they are not generated through global inferences, but are conventionally associated with certain linguistic elements, and on the other hand differ from conventional implicatures in that they are cancelable. This seems to be an adequate description of the above hypothesized interpretation of a word like *Turk*. As, like conventional implicatures, these default implicatures are conventionally tied to particular expressions, Recanati (2003) argues that they can be considered semantic rather than pragmatic. This is in line with what could be called ‘rich’ approaches of the lexicon in which words contribute rich semantic representations to the process of interpretation and in which meaning aspects do not necessarily have to be non-defeasible to be part of the lexical, semantic meaning of words (e.g. Hogeweg & Vicente 2020). Zeevat et al. (2017), for example, describe lexical meanings as distributions over bundles of (moderately universal) semantic features. These feature sets can be specified by a number of operations, such that some features may be absolute, while others are e.g. mutually exclusive. A similar approach might also be applicable to the loss of descriptive content, as seems to (partly) be the case for *flikker*. In this case, it is the descriptive reading that becomes less conventional and more dependent on the context.

Independently of the precise analysis or categorization of the various terms, my main point is that it may not be that easy to make a strict distinction between terms that do and terms that do not ‘have the expressive function hardwired into their lexicon’ and between terms that do or do not (still) have descriptive content. The three stages of pragmaticalization described by Gutzmann & Davis (2015) may therefore best be seen as points in a continuous process. This does not mean that there is no distinguishable class of words that fit all the criteria commonly ascribed to slurs, but that such terms are part of a continuum of (more or less) pejorative terms for (groups of) people.

A gradual view on conventionality in the lexical domain would also pave the way for a less fundamental distinction between expressive and social meaning. Henderson & McCready (2019) argue that unlike slurs, dog whistles (i.e., the phenomenon whereby a speaker sends one message to an outgroup and a second, often controversial message to an ingroup) are not properly analyzed in terms of conventional implicature. An important argument against a conventional implicature account is that the ingroup message can be plausibly denied by a speaker. Henderson & McCready (2019) instead analyze dog whistles as conveying social meaning; they signal a certain persona. Interestingly, Nunberg (2018) provides a similar analysis for slurs. Nunberg (2018) argues against an analysis of slurs in terms of conventional implicature and claims that slurs get their offensive nature through speakers’ exploitation of sociolinguistic variation and their affiliation with speakers who ‘own the word’. The question arises whether there is a need for such a strict

opposition between conventional implicature and social meaning. Indeed, Eckert (2019: 769) argues for an integrated view on meaning in which “sociolinguistic variation contributes a purely performative, subtle, speaker-indexical resource, ranging from social category membership to momentary affective states.” As such, social meaning and expressive meaning conveyed through conventional implicature are not fundamentally different types of meaning but they form “a cline of ‘interiority’ from variables that index public social facts about the speaker to more internal, personal affective states.” (Eckert 2019: 751).

Based on the observations with respect to Dutch slurs and related words presented in this paper, I suggest that slurs are part of a larger collection of pejorative or expressive terms for (groups of) people, and that individual terms within this collection can show mixtures and degrees of descriptive, expressive and social meaning. Instead of aiming for ‘a theory of slurs’, it might therefore be more useful to investigate individual or subclasses of terms to find out the range of combinations of meaning they can express and the way they are interdependent and interact. A good example of such an analysis is Burnett’s (2020) analysis of the terms *d**** and *lesbian*. Burnett shows that *d**** and *lesbian* are not descriptively equivalent but that they are associated with different personae; *lesbian* is associated with a mainstream persona, and *d**** is associated with what Burnett describes as an anti-mainstream persona. The appreciation for the different personae differs between people having different ideologies. The interpretation of the terms is then determined by an interaction between the meaning of the term in the listener’s own ideology and their expectations about the speaker’s ideological structure. Burnett (2020: 57) suggests the possibility that a similar analysis would be helpful for other terms that have been analyzed “under the umbrella of the term ‘slur’”, such as *f*** or *slut*, but also for ‘non-derogatory’ terms for women as a *good woman* or a *nice young lady*.

4 Conclusions

There is a lot of debate and little consensus in the philosophical, linguistic and (formal) semantic literature about slurs. Typically, authors argue against an analysis in terms of one category of meaning (presupposition, conversational implicature, conventional implicature, social meaning, etc.) and for the analysis of slurs in terms of another. The discussion on slurs is furthermore characterized by its highly theoretical nature. Based on the observations with respect to Dutch slurs and related words presented in this paper, I suggest that the debate would benefit from a broader and more empirical approach to pejorative terms, as their actual use and interpretation can show a mixture of (degrees of) descriptive, expressive and social meaning. This does not mean that ‘true slurs’ do not exist or are not unambiguously identifiable, but they may be rarer than assumed once one considers their actual use and interpretation and, furthermore, that other words, although not slurs themselves, are interesting and relevant in light of the discussion of pejorative and expressive language.

Acknowledgements

I am grateful to Wessel Stoop for his help in collecting the data. I would also like to thank Michelle Suijkerbuijk, Caya van Dijk, and Elsa Opheij for their help with the annotation of the data and/or finalizing the text. Furthermore, I would like to thank the audience of the *Grote Taaldag* and the *22nd Szklarska Poręba Workshop* for their useful comments. Finally, I thank three anonymous reviewers and the editors of *Glossa* for their very useful comments and advice.

Competing interests

The author has no competing interests to declare.

References

- Bartlett, Jamie & Reffin, Jeremy & Rumball, Noelle & Williamson, Sarah. 2014. *Anti-Social Media*. London: Demos.
- Bolinger, Renée Jorgensen. 2017. The pragmatics of slurs. *Noûs* 51(3). 439–462. DOI: <https://doi.org/10.1111/nous.12090>
- Burnett, Heather. 2020. A persona-based semantics for slurs. *Grazer Philosophische Studien* 97(1). 31–62. DOI: <https://doi.org/10.1163/18756735-09701004>
- Camp, Elisabeth. 2013. Slurring perspectives. *Analytic Philosophy* 54(3). 330–349. DOI: <https://doi.org/10.1111/phib.12022>
- Camp, Elisabeth. 2018. A dual act analysis of slurs. In David Sosa (Ed.), *Bad words*, 29–59. Oxford: Oxford University Press. DOI: <https://doi.org/10.1093/oso/9780198758655.003.0003>
- Carnaghi, Andrea & Maass, Anne. 2007. In-group and out-group perspectives in the use of derogatory group labels: Gay versus fag. *Journal of Language and Social Psychology* 26. 142–156. DOI: <https://doi.org/10.1177/0261927X07300077>
- Carnaghi, Andrea & Maass, Anne & Gresta, Sara & Bianchi, Mauro & Cadinu, Mara & Arcuri, Luciano. 2008. Nomina sunt omina: On the inductive potential of nouns and adjectives in person perception. *Journal of Personality and Social Psychology* 94(5). 839–859. DOI: <https://doi.org/10.1037/0022-3514.94.5.839>
- Croom, Adam M. 2011. Slurs. *Language Sciences* 33(3). 343–358. DOI: <https://doi.org/10.1016/j.langsci.2010.11.005>
- Croom, Adam M. 2013. How to do things with slurs: Studies in the way of derogatory words. *Language & Communication* 33(3). 177–204. DOI: <https://doi.org/10.1016/j.langcom.2013.03.008>
- Eckert, Penelope. 2019. The limits of meaning: Social indexicality, variation, and the cline of interiority. *Language* 95(4). 751–776. DOI: <https://doi.org/10.1353/lan.2019.0072>
- Fasoli, Fabio & Carnaghi, Andrea & Paladino, Maria Paola. 2015. Social acceptability of sexist derogatory and sexist objectifying slurs across contexts. *Language Sciences* 52. 98–107. DOI: <https://doi.org/10.1016/j.langsci.2015.03.003>

- Foolen, Ad. 1997. The expressive function of language: Towards a cognitive semantic approach. In Niemeier, Susanne & Dirven, René (eds.), *The Language of Emotions: Conceptualization, Expression, and Theoretical Foundation*, 15–31. Amsterdam: John Benjamins. DOI: <https://doi.org/10.1075/z.85.04foo>
- Grice, H. P. 1975 . Logic and conversation. In Cole, Peter & Morgan, Jerry L. (eds.), *Syntax and Semantics 3: Speech Acts*, 41–58. New York: Academic. DOI: https://doi.org/10.1163/9789004368811_003
- Grice, Paul. 1989. *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Gutzmann, Daniel. 2015. *Use-conditional meaning: Studies in multidimensional semantics*. Oxford: Oxford University Press. DOI: <https://doi.org/10.1093/acprof:oso/9780198723820.001.0001>
- Gutzmann, Daniel. 2019. *The Grammar of Expressivity*. Oxford: Oxford University Press. DOI: <https://doi.org/10.1093/oso/9780198812128.001.0001>
- Gutzmann, Daniel & Davis, Christopher. 2015. Use-conditional meaning and the semantics of pragmaticalization. *Proceedings of Sinn and Bedeutung* 19. 197–213.
- Henderson, Robert & McCready, Elin. 2019. Dogwhistles and the at-Issue/non-at-issue distinction. In Gutzmann, Daniel & Turgay, Katharina (eds.), *Secondary Content*, 222–245. Leiden, The Netherlands: Brill. DOI: https://doi.org/10.1163/9789004393127_010
- Hogeweg, Lotte & Neuleman, Monique. 2022. The hurtfulness of slurs, nouns and adjectives as group labels. *Linguistics in the Netherlands* 39(1). 72–87. DOI: <https://doi.org/10.1075/avt.00062.hog>
- Hogeweg, Lotte & Vicente, Agustin. 2020. On the nature of the lexicon: The status of rich lexical meanings. *Journal of Linguistics* 56(4). 865–891. DOI: <https://doi.org/10.1017/S0022226720000316>
- Hom, Christopher. 2008. The semantics of racial epithets. *Journal of Philosophy* 105. 416–440. <https://www.jstor.org/stable/20620116>. DOI: <https://doi.org/10.5840/jphil2008105834>
- Hornsby, Jennifer. 2001. Meaning and uselessness: How to think about derogatory words. *Midwest Studies In Philosophy* 25(1). 128–141. DOI: <https://doi.org/10.1111/1475-4975.00042>
- Jeshion, Robin. 2016. Slur creation, bigotry formation: the power of expressivism. *Phenomenology and Mind* 11. 130–139. DOI: https://doi.org/10.13128/Phe_Mi-20113
- Kaplan, David. 1999. *The meaning of ouch and oops: Explorations in the theory of meaning as use*. Los Angeles: University of California, Los Angeles manuscript.
- McCready, Elin. 2010. Varieties of conventional implicature. *Semantics and Pragmatics* 3(8), 1–57. DOI: <https://doi.org/10.3765/sp.3.8>
- Nunberg, Geoff. 2018. The social life of slurs. In Fogal, Daniel & Harris, Daniel & Moss, Matt (eds.), *New Work on Speech Acts*, 237–95. Oxford: Oxford University Press. DOI: <https://doi.org/10.1093/oso/9780198738831.003.0010>
- O’Dea, Conor J. & Miller, Stuart S. & Andres, Emma B. & Ray, Madelyn H. & Till, Derrick F. & Saucier, Donald A. 2015. Out of bounds: Factors affecting the perceived offensiveness of racial slurs. *Language Sciences* 52. 155–164. DOI: <https://doi.org/10.1016/j.langsci.2014.09.005>

- Potts, Christopher. 2005. *The logic of conventional implicatures*. Oxford: Oxford University Press. DOI: <https://doi.org/10.1093/acprof:oso/9780199273829.001.0001>
- Recanati, François. 2003. Embedded implicatures. *Philosophical Perspectives* 17. 299–332. <https://www.jstor.org/stable/3840877>. DOI: <https://doi.org/10.1111/j.1520-8583.2003.00012.x>
- Schmid, Hans-Jörg. 2015. A blueprint of the entrenchment-and-conventionalization model. *Yearbook of the German Cognitive Linguistics Association* 3(1). 3–25. DOI: <https://doi.org/10.1515/gcla-2015-0002>
- Sternau, Marit & Ariel, Mira & Giora, Rachel & Fein, Ofer. 2017. Deniability and explicatures. In Giora, Rachel & Haugh, Michael (eds.), *Doing Pragmatics Interculturally*, 97–120. Berlin/Boston, MA: Walter de Gruyter. DOI: <https://doi.org/10.1515/9783110546095>
- Townsend, Leanne & Wallace, Claire. 2016. Social media research: A guide to ethics. University of Aberdeen, available at: <https://www.bolton.ac.uk/assets/Uploads/Social-media-ethics-study-Aberdeen-2018.pdf>
- Williamson, Timothy. 2009. Reference, inference and the semantics of pejoratives. In Almog, Joseph & Leonardi, Paolo (eds.), *The Philosophy of David Kaplan*, 137–158. Oxford: Oxford University Press. DOI: <https://doi.org/10.1093/acprof:oso/9780195367881.003.0009>
- Zeevat, Henk & Grimm, Scott & Hogeweg, Lotte & Lestrade, Sander & Smith, E. Allyn. 2017. Representing the lexicon: Identifying meaning in use via overspecification. In Balogh, Kata & Petersen, Wiebke (eds.), *Bridging Formal and Conceptual Semantics: Selected papers of BRIDGE-14*, 153–187. Düsseldorf: düsseldorf university press. DOI: <https://doi.org/10.1515/9783110720327-007>

