



Slurs and Redundancy

Y. Sandy Berkovski¹ 

Received: 29 September 2021 / Revised: 4 January 2022 / Accepted: 2 February 2022 /
Published online: 6 April 2022
© The Author(s), under exclusive licence to Springer Nature B.V. 2022

Abstract

According to nearly all theorists writing on the subject, a certain derogatory content is regularly and systematically communicated by slurs. So united, the theorists disagree sharply on the elements of this content, on its provenance, and on its mechanism. I argue that the basic premiss of all these views, that there is any such derogatory content conveyed with the use of slurs, is highly dubious.

Keywords Slurs · Derogation · Pragmatics · Semantics

1 Introduction

According to many, nearly all theorists writing about slurs, there is a certain derogatory content *D* slurs carry on their sleeves. *D* is associated with slurs in a lawlike manner, so that under normal circumstances, when a speaker uses a slur, she will communicate *D*. The shared assumption was stated clearly in Christopher Hom's writings:¹

One of the main distinguishing features of racial epithets is their capacity to derogate their intended targets in deep and explosive ways. The derogatory force for any epithet is *independent* of the attitudes of any of its particular speakers. For example, uses of 'chink' carry the same derogatory force no matter how racist or nonracist the particular speaker is towards Chinese people. (Hom, 2008, 426, italics added)

Derogation is an *objective* feature of the semantic contents of pejorative terms. Derogation is the result of the actual predication, or application, of a slur or pejorative term to its intended target group. (Hom, 2012, 397, italics added)

¹Another early general statement to the same effect: Hornsby (2002: 134–135). A recent statement: Torrenco (2020:1626).

✉ Y. Sandy Berkovski
sandy.berkovski@gmail.com

¹ Department of Philosophy, Bilkent University, Ankara 06800, Turkey

Theorists disagree on the precise nature of derogation and on the linguistic mechanism that attaches it to slurs. But they agree that a derogatory ‘content’, ‘force’, or ‘attitude’ *is* attached to slurs.² Hence they must also hold that, at least under average circumstances and in simple subject-predicate constructions, this derogation is contributed to the content of the whole utterance.

I believe that this view implausibly associates the use of slurs with regular pragmatic infelicities. Specifically, in Section 2 I argue that, on the objectivist view just mentioned, *D* becomes redundant: the association of *D* with slurs yields a violation of the maxim of quantity. This is so at least when the same slur is used repeatedly in the conversation. In Section 3 I respond to several possible objections. Some might claim that violations of the maxim of quantity are, and should be, tolerated. I reply that the occasions where the maxim of quantity is felicitously violated are irrelevant for the case of slurs. Others might claim that slurs are rarely, if ever, used repeatedly in the same conversation. Using some statistical data, I reply that the typical use of slurs is in the conversations among bigots where we expect to find the same slur used repeatedly. In Section 4 I assess the impact of the redundancy argument on the debate over slurs. I argue that ‘prohibitionism’ is one view immune to the redundancy argument. If only for this reason it deserves another look. I finish by briefly addressing a few familiar objections to prohibitionism.

2 The Argument

Let ‘implicaturism’ be the view that slurs conventionally implicate a certain content *D*. It is defended in Williamson (2009) and Whiting (2013). Williamson gives the examples:

- (1) a. Lessing was a German.
b. Lessing was a Boche.

Both utterances have the same truth-conditional content. But only the latter conventionally implicates that Germans are cruel. Whiting’s view is slightly different. There is no truth-conditional content conventionally implicated by (1b). Rather, the utterance is ‘expressive’ of a non-cognitive attitude. I take it that on this view, the utterance is supposed to communicate the speaker’s bigoted attitude toward the Germans.

Implicaturism is plausible when we think of a speaker using a slur for the first time in a conversation. If in the course of a sterile academic debate about Lessing you burst out with (1b), no incongruities are detected. But suppose that you use the slur repeatedly in the conversation, matter of factly, and your interlocutors follow the suit. Then we have a problem. Consider this accretion of redundant content that results in infelicities:

- (2) a. (A:) — Lessing was a Boche.

²In the useful terminology of Torrenco (2020:1618) this consensus cuts across the ‘strict content’ and ‘broad content’ approaches to slurs. Moreover, some objectivists adopt a pragmatic explanation of derogation (see below).

- b. (B:) — And d’Holbach was probably a Boche.
- c. (A:) — Not sure. But Wieland was definitely a Boche.
- d. (B:) — Yes, and Hamann was a German, too.

It was informative for *A* to convey a derogatory content with his first utterance. But if the *same* implicature attaches to every use of slur, then repeating the same content all over again would be foolish. If *A* has already voiced (implicated) his opinion or attitude about the Germans in (2a), why would he go on and voice it again?

Given the redundancy, should the speaker stop using the slur after its first introduction? She certainly may, but the risk is she’d be misunderstood. In (2d), where the speaker is doing exactly that, *B* may be interpreted as expressing different attitudes to Hamann and d’Holbach. Perhaps, the reasoning goes, *B* used the conventional implicature in (2b) to establish her own manner of referring to Germans. Coming to Hamann, she violates her newly established manner, and so in this sense flouts the maxim of manner. In other words, *B*—or more accurately, *B*’s utterance—first conventionally implicates that Germans are cruel, and that Lessing is one of those Germans. *B* then conversationally implicates that Hamann is not one of the cruel Germans.

In general, however, we are not entitled to make this inference. Consider an example more realistic than the artificial (2): a scene from a motion picture *The Inglorious Basterds*. Aldo Raine (played by Brad Pitt) and his ‘basterds’ have just captured a group of German soldiers. In the course of his interrogation of a German soldier, Lt. Raine is heard uttering these sentences, all within a span of a few minutes:

- (3) a. Send that Kraut sarge over.
- b. Besides you, we know there’s another kraut patrol fuckin’ here somewhere.
- c. I need to know about Germans hiding in trees.
- d. Got us a German here who wants to die for country.

We observe two characteristics in this sort of talk:

High-frequency usage A slur may felicitously be used in multiple utterances by the same speaker, in the same conversation.

Alternating usage A slur may routinely be substituted for its neutral correlate, and *vice versa*, in multiple utterances in the same conversation.

Raine is slipping in and out of the pejorative use without a second thought, and we follow him easily. His initial choice of ‘Kraut’—initial for us, but probably not for his audience of the basterds—perhaps suggests to us something about Raine’s beliefs or his character. But on the second occasion, it’s unlikely that anything of this kind is suggested or implicated. Moreover, the fact that, without any warning, Raine switches to the neutral usage is equally insignificant. In particular, we are not to infer that Raine changed his view of the German patrol or of the captured soldier, or of the Germans generally.³

³To prevent a misunderstanding: I’m not claiming that the film director Quentin Tarantino has any special expertise in bigoted use, nor that any such expertise exists in the first place (Nunberg, 2018, 283). I am rather claiming that Tarantino has portrayed the repeated use fairly convincingly, and that we find this use entirely fluent and unproblematic.

Somewhat unexpectedly, analogous considerations tell against pragmatic accounts. According to Nunberg (2018), when a speaker chooses to use a slur, he conveys his affiliation with the relevant bigoted group. This lies at the heart of a ‘ventriloquistic implicature’ associated with slurs:

In a particular context, a speaker pointedly disregards the lexical convention of the group whose norms prescribe the default way of referring to A and refers to A instead via the distinct convention of another group that is known to have distinct and heterodox attitudes about A, so as to signal his affiliation with the group and its point of view. [...] [T]he speaker *always* has an ulterior reason for using them, over and above the proposition he asserts. (Nunberg, 2018, 267, 269, italics added)

Ventriloquistic implicature is a generalised conversational implicature. Since slurs are always marked alternatives to the conversational default, ventriloquistic implicatures are regularly generated by slurs. This idea may explain the initial use possibly designed to establish one’s bigoted credentials. But once again, the High-frequency and Alternating usage of slurs are puzzling. People in the audience should be perplexed about both the repeated presence and occasional inexplicable absence of ventriloquistic implicatures—whereas in reality they are not.

Consider now the view in Bolinger (2017). Its main idea is that the use of slurs is a ‘signal’ of a certain cluster ϕ of attitudes:

Speakers competent with a language have knowledge not only of lexical items and grammar, but also a set of co-occurrence expectations that encode the social norms and conventions concerning the use of various terms and ways of speaking. At their most general, such expectations are of the form ‘the behavior α characteristically signals ϕ ,’ where ϕ ranges over some associated information. [...] Use of a slur α rather than a neutral alternative β is a defeasible indicator that the speaker endorses ϕ , and can be undercut by speaker ignorance or a forced choice. (Bolinger, 2017, 447–448)

Whether this account is threatened by the redundancy argument turns on the exact theoretical role of ‘signals’ and ‘indicators’. An informal talk of slurs as ‘signals’ is, of course, harmless. Informally, signals ‘transmit’, ‘convey’, ‘communicate’ information, and slurs, like any linguistic terms, can do the same.

Yet Bolinger uses the term in its technical sense. She locates slurs explicitly within the signaling framework of Skyrms (2010). But then, are slurs part of the signaling game between senders (slurs’ users) and receivers? If they are, who are the receivers, exactly? Bolinger doesn’t say, but since the explanandum is the offence generated by slurs, we speculate that the receivers are the offended party, the Victims. But this won’t do: for a signaling game to be in equilibrium, both senders and receivers must have sufficient interest in receiving the information transmitted by the signal. It may be in the Victims’ interest to learn who the Bigots are so as to avoid them, but this would be offset by emotional suffering. It is similarly unclear what payoff the Bigots would derive from broadcasting their bigoted attitudes to the world, the Victims included. More plausibly we should say that senders and receivers are fellow bigots who use slurs to identify their ilk and to identify themselves to them, as suggested by

Nunberg. But this communication will be vitiated exactly by the High-frequency and Alternating usages.

Generally, supposing that slurs *are* part of the signaling game between some senders and some receivers, then they are used for communication: slurs are used to communicate the attitudes ϕ . But this is just another way of saying that the attitudes ϕ are part of the slurs' meaning. Then the redundancy argument goes through: the repeated use of slurs must result in the accretion of redundant content, just as we saw before.

There is, however, another view that could be read into Bolinger's quote above (also into some of Nunberg's remarks): slurs are mere attributes of behaviour, like accent, heart rate, or perspiration.⁴ When Henry Kissinger speaks English with a German accent, he does not 'convey' or 'imply' that he is a native German speaker. Nevertheless the hearer may infer a number of things. A rational conclusion (if not always a 'warranted' one, in Bolinger's terminology) would be that Kissinger is a native German speaker. The speaker no longer conveys any content or attitude. Instead, we look at the exchange entirely from the hearer's perspective. Under the normal circumstances, your use of 'kike' raises my credence in the proposition that you have derogatory attitudes toward the Jews. You might not wish to communicate any past or current attitudes toward the Jews, and your utterance in itself doesn't communicate them either. Still, given the fact of your use, it would be rational for me to assign a higher probability to the belief that you are a bigot.

This natural signaling view would be immune to the redundancy argument. Kissinger's accent is not redundant simply because Kissinger spoke a lot, though it would be redundant for you to use it repeatedly to draw the same conclusions about Kissinger. The speaker's utterances aren't infelicitous: the only redundant thing may be the hearer's inferences. So should we understand slurs as natural signals, mere attributes of behaviour? There are clear signs that Bolinger does not intend this interpretation. Slurs are means of *communication*. The user of a slur is said to be 'aware' of the signaling relation, to 'choose' to use the slur, to 'signal' that he 'endorses' the derogatory attitudes (448–450). Bolinger's account, as presented, is vulnerable to the redundancy argument.

To touch briefly on an approach in a different tradition, consider a recent speech-theoretic account in Liu (2021) where utterances containing slurs are interpreted as having the illocutionary force of derogation. Derogatory speech acts are grouped with declaratives. Their essential condition, according to Liu, is the intention to enforce discriminatory norms against the target. Its propositional content is 'persons or groups of people'. How exactly this is spelled out further doesn't really matter. In particular, it doesn't matter what exactly, according to this view, is supposed to be communicated by an utterance like 'A is a kraut.' It's enough only to concede that, e.g., Aldo Raine's utterances *are* saddled with a certain invariant content. Then the redundancy argument goes through.

But isn't there a simple tweak to apply to the accounts we have examined that would make them invincible to the redundancy argument? Rather than saying that

⁴A similar ambiguity, I think, vitiates the perspectivalist 'semantic' account in Camp (2013), Camp (2018).

an *invariant* derogatory content is associated with a sentence like ‘A is a kraut’ you might say that the derogatory content is variable and tied to A. So that sentence would be analysed along the lines of ‘A is German and despicable.’ It won’t communicate anything about the Germans generally. In fact this idea is already favoured by some theorists that integrate the derogatory content straight into the semantic content. Consider Christopher Hom’s ‘semantic externalism’:

Basically, to call someone a D is to *say* that they ought to be subject to discriminatory practices for having negative, stereotypical properties because of being an N (Hom, 2010, 180, italics added),⁵

where ‘D’ and ‘N’ are a slur and its neutral counterpart, respectively. On this view, the predicate ‘ ξ is a kraut’ designates a complex property ‘ ξ ought to be subject to a discriminatory treatment *X* because of certain negative properties *Y* that ξ has because of being German’. Similarly, according to Robin Jeshion’s version of expressivism, the predicate ‘ ξ is a kraut’ semantically ‘encodes’ the expression of contempt toward the particular individual qua German.⁶

Yet this tweak can’t disable the redundancy threat. The reasoning is analogous to the one we used with (2) and (3). Imagine the following conversation among the ‘basterds’:

- (4) a. (Raine:) — I need to know about those krauts hiding in trees.
 b. (Wicki:) — They aren’t there.
 c. (Ulmer:) — I wonder where the krauts went.
 d. (Raine:) — These Germans are sly.

As the conversation progresses, the High-frequency and Alternating usages become more prominent. Different individuals are sometimes described as ‘German’, sometimes as ‘kraut’, with no regular pattern. Then you are apt to conclude that ‘kraut’ and ‘German’ are in fact synonyms, and that ‘kraut’ isn’t meant to communicate anything extra about the individuals in question.

In the case of expressivism, it is useful to present this reasoning by considering gestures. Without getting into their semantics, it’s uncontroversial that certain gestures, such as showing the middle finger, may express a certain derogatory attitude *d*.⁷ What *d* exactly is, contempt or something else, doesn’t matter. We only have to imagine that Aldo Raine has the attitude *d* toward the Germans. Then, whilst referring to the Germans with a neutral term, he could indicate his attitude by showing the middle finger. The first occurrence of the gesture could be interesting and informative. But plainly, if Raine persists with this gesture, we should begin to worry about his mental condition. The same should go for the terms like ‘Kraut’ or ‘Boche’. Compulsively expressing attitudes through utterances is no less odd than repeating the same gesture again and again.

⁵See also Hom (2008:431).

⁶See Jeshion (2018:82–83, 100). Note that expression of attitudes is subject to the maxim of quantity. To obey this maxim we shouldn’t make our utterances more ‘informative’ than necessary, and to express an attitude *is* to communicate information. Maxim of quantity: Grice (1989:26).

⁷Gestures and slurs: Hornsby (2002:140).

I conclude that there is no simple variability tweak to do away with redundancy. But should redundancy be recognised as a threat in the first place? This is the subject of Section 3.

3 Objections

In this section I look at the possible objections to the cogency or significance of the redundancy argument.

(I) The first objection might come from the implicaturist. Don't we have the same redundancy with conventional implicatures generally, and wouldn't they similarly result in infelicities? For example, we often use 'even' in the same conversation. Even if its use should be judged redundant, we don't seem to mind. Same for slurs. The purported redundancy is not a threat to the coherence of conversation. It can't be a reason to dismiss the theories that entail it.

I reply that in the redundancy argument we have assumed that the *same* content was produced at the different stages of the conversation. Familiar conventional implicatures are not threatened by redundancies, so far they are tied to the at-issue content. Consider 'even' in the following exchange:⁸

- (5) a. (A:) — Even Ken knows it's unethical. (Conventional implicature: Ken is the least likely person of the salient group to know it's unethical.)
 b. (B:) — Well, even Ann knows it's unethical. (CI: Ann is the least likely person ...)
 c. (A:) — After all, even Jo knows it's unethical! (CI: Jo is the least likely person ...)

No redundancies here, since a *new* content is implicated with each utterance. It is different with Williamson's account, for example: there, the same content 'Germans are cruel' would be implicated with each use of 'Boche'.

(II) A more general objection on the same theme is to challenge directly the status of the maxim of quantity. Multiple studies in linguistics and psychology show that conversation participants routinely violate the maxim. Let me divide these studies into three groups.

Studies in the first group concern repetitions.⁹ In longer strips of discourse participants often repeat intonations, phonemes, sub-sentential expressions, and then also sentential expressions. With the last one in particular this results in the repetition of content, and thus in a violation of the quantity maxim. These repetitions are interesting, but hardly mysterious. They grease the wheels of social interactions by emphasising the significance of different subjects, strengthening the emotional bonds between the participants, confirming hints, querying information, or ratifying it (A: 'He is dead!' B: 'Yes, he is dead!'). Or they may be the standard syntactic forms of

⁸See Horn (2004:3). The reader is invited to test other instances on the list in Potts (2015:188).

⁹Concise summary: Brown (2000). Further details: Tannen (2007), Schegloff (1996).

certain registers like poetry or oratory ('We shall fight on the beaches, we shall fight on the landing grounds, we shall fight in the fields and in the streets etc.').

Studies in the second group concern referential overspecifications.¹⁰ Speakers would often refer to an obviously salient entity with redundant modifiers. If there is only one bowl and several plates lying on the table, I might still tell you, 'Pass me the red bowl, please.' The request is felicitous, though the modifier 'red' is redundant, and its use conflicts with the maxim of quantity. One explanation is that overspecifications reduce the cognitive load of the hearer. An alternative explanation is that they in fact reduce the cognitive load of the speaker, i.e. the costs of his speech-production process.

Studies in the third group concern the special circumstances of strategic interactions where the maxim of quantity is 'deactivated'.¹¹ For example, suppose there is a game show where money is put in five boxes out of hundred. Then an utterance 'There is money in box 20 or 25' made by a game show host is not judged infelicitous. This is so, although the host is supposed to know exactly where money is. Nor do we take this utterance to indicate that the host is uncertain about where money is put. Both of these judgements are explained by the fact that, in the given context of competition or conflict, the speaker is not expected to obey the maxim of quantity.

Therefore, since the maxim of quantity is by no means sacrosanct, the current objection runs, no notable infelicity is produced by the use of slurs even if slurs are endowed with the invariant derogatory content. I reply that the cases aren't similar. Clearly the third type of situation is irrelevant here. The use of slurs is not normally supposed to violate the maxim of quantity because of the unique requirements of context. Nor does it fit well the second type of situation. The use of slurs is not a way to make reference production more effective. Nor is its putative redundancy limited to the redundancy of modifiers.

The objection may seem to be on a firmer ground in the first type of situation. The violation of the maxim of quantity might be justified by the same sort of repetitions. The redundancy of the derogatory content, the opponent would argue, is occasioned by the need to project emotion, solidarity (talking to fellow bigots), or hostility (talking to victims). Bigots repeatedly use slurs to create a certain atmosphere in the conversation and to re-affirm their bigoted commitments. But there are several reasons to doubt this idea, as follows.

(a) First, the recurrent use of slurs in a conversation like (3) may and does persist even when all the purposes mentioned above have conceivably been accomplished. At some fairly early point in the conversation (3) all the participants will be very clear about who believes or feels what about the Germans. But that wouldn't prevent them from using 'kraut' perfectly felicitously. (b) Second, our earlier observation stands: at some early point in the conversation the derogatory generic content $D = \text{'Germans are } F\text{'}$ isn't processed by the participants (even assuming that it was

¹⁰Review of the debate: Davies and Arnold (2019). Further details: e.g., Rubio-Fernandez (2019), Engelhardt et al. (2006).

¹¹See Fox (2014).

processed at the start of the conversation). Aldo Raine’s audience simply don’t interpret his utterances as conveying that content. Therefore, if it’s real at all, *D* is idle and unlikely to be used for any purpose whatever. (c) Thirdly, and relatedly, a repetition like ‘He is dead!’ is a repetition of what was said before. Already at a pre-theoretic level, ignoring the maxim of quantity, it is somewhat unusual and calls for an explanation. But Aldo Raine’s use of ‘kraut’ doesn’t even seem like a repetition of the same content to be explained. Instead, it seems more like a repeated *identification* of individuals and groups with the same lexical item—a completely banal phenomenon. Of course the theories of slurs may interpret his utterances as repetitions of content, but that’s just the reason to doubt them. Intuitively, it’s implausible to think of them as repetitions of content in the first place.

(III) The third objection I want to look at is that repeated use is perhaps too aberrant to merit a radical rethink of derogatory content. The redundancy argument may be valid, but has marginal theoretical interest, so far as it applies to a rather marginal phenomenon. In other words, this objection rests on the assumption *R*, that repeated use is sufficiently rare to merit any scrutiny.

So *is* repeated use rare? You might reason as follows. If we take the ‘out-group’ use by bigots against their targets (the Victims) as the primary provenance of slurs, then conceivably slurs are met with indignation, protests, possibly violence. We could interpret them as ‘fighting words’ and adapt the legal reasoning in the classical *Chaplinsky* case where Justice Murphy defined ‘fighting words’ as those that ‘tend to incite an immediate breach of peace’.¹² Just as in the *Chaplinsky* the court judged the expressions ‘racketeer’ and ‘fascist’ fighting words, we could extend the same judgement to every slur. Thus *R* would be vindicated: slurs can’t occur in any given interaction too frequently, since such an interaction would be cut short by disturbances.

But the assumption *R* is likely false. To see why, consider this key observation due to Geoff Nunberg (logically independent of his theory of slurs’ meaning that we challenged earlier):

[T]he focus on the offensiveness of slurs tends to obscure what is usually their primary *raison d’être*. We should bear in mind that vast majority of the uses of these words occur among the members of the group they belong to, out of earshot of any of the people they target. Indeed, a community may have a slur for a group of people that its members have no expectation of ever encountering. (Nunberg, 2018, 253)

The primary provenance of slurs, Nunberg suggests, is the ‘in-group’ use among the bigots. A quick armchair investigation confirms Nunberg’s hypothesis: the directly offensive use of antisemitic slurs is, at present, exceedingly rare, while the number of antisemitic incidents keeps rising. In the ADL tracker of these incidents in the United States in 2019–2020, by my counting only three out of several thousands of incidents explicitly mentioned the use of slurs (there were in total 2,107 antisemitic incidents

¹²As Arkes (1974) notes, Murphy’s definition was later adopted nearly verbatim in several influential rulings.

Table 1 Nexis Uni data

	2019	1975–2001	?–2001 (Kennedy)
N-word	310	3,333	4,219
Kike	5	56	84
Wetback	22	164	50

reported in 2019 alone).¹³ In the annual 2019 summary only one such incident is mentioned, and even then it probably wasn't meant to be addressed to the targeted group.

Amateur and unscientific as it may be, this statistic points to two conclusions. Although there is no shortage of them in the English vernacular (unlike in German, for example), slurs are very rarely used in addressing the targeted groups in such a way as to merit reporting. Given that derogation itself is a rare phenomenon, this also means that slurs are even more seldom used in regular interactions familiar to non-bigots. Slurs are exotic. Joined with a reasonable premiss that they are not a dead artefact of the language and *are* actually used, we conclude that they must be used in other contexts, quite likely among the bigots themselves. This is in line with Nunberg's hypothesis.

Our amateur statistic may be challenged by citing another (amateur) statistic in Kennedy (2002:26) that draws on the LexisNexis database and is purported to show the widespread use of slurs by bigots directed at their targets. The data contain the number of references to the relevant terms in legal cases across the United States. I was unable to replicate Kennedy's results in Nexis Uni, and instead in Table 1 I give the comparable data for 1975–2001 and for the single year 2019.¹⁴

These data contain a lot of noise. For example, among the five cases in 2019 where 'kike' was mentioned in court documents only one was a case of actual derogation where the defendant used the slur. Even superficially, it is clear that similar noise, though not of the same magnitude, is in other cases. What's more, many of the cases contain reports of in-group use where, e.g., the plaintiff complains about the racist environment in which slurs were used in conversations among bigoted co-workers. Supposing, however, we take the data at face value and treat all these cases as instances of out-group use, it still doesn't follow that slurs are 'frequently' directed at victims, or that they are so used more frequently than among the bigots themselves. To know any of this we need to know the number of bigots and the frequency of slur use among them. This, of course, is difficult to measure, if possible at all.

We are left with anecdotal and indirect evidence. As a member of a discriminated minority <DETAILS REDACTED>, I can report that the relevant slurs were directed at me perhaps at most three or four times in the course of my whole life (forty-nine years). Randall Kennedy himself narrates movingly his experiences of discrimination, yet the actual uses of the N-word that he mentions are few. This is not to

¹³ADL tracker: <https://bit.ly/2M84W7R> ADL 2019 audit: <https://www.adl.org/audit2019>

¹⁴Nexis Uni: <http://www.nexisuni.com/>

dispute that every exposure to such use was painful. In fact, those experiences could be particularly painful and memorable precisely because they were rare. Turning to indirect evidence, in Waugh's fairly long *Officers and Gentlemen* set in Italy the slur 'wop' is mentioned only once, and predictably in an exchange between two English characters. In Hemingway's partly semi-documentary *In Our Time* 'wop' is freely used among non-Italian speakers, but never addressed to Italians, though some of the action is set in Italy as well.

It may be objected that, whatever the case with other slurs, the N-word is a clear outlier. Even in our toy statistic the frequency of its use is staggering when compared to other slurs. But as already noted, the numbers may tell as much about the extent of bigotry in the population, as about the extent of out-group use. Again, in Hemingway's *Sun Also Rises* and *To Have and To Have Not* we have multiple uses of the N-word, but nearly exclusively in the conversations among whites.

I don't mean to suggest that we have here a clear proof of Nunberg's hypothesis. I only suggest that Nunberg's hypothesis, amateur statistic, and anecdotal evidence are mutually supportive. Thus the hypothesis is sufficiently plausible to be given at least some weight in our theorising. If we take Nunberg's hypothesis on board, we can summarily dismiss the assumption *R*. Slur use need not generally be rare at all; in fact, it's likely to be frequent. In the in-group settings bigots are likely to use slurs as frequently as any other words. Hence the repeated use and the resulting violations of the maxim of quantity aren't so aberrant, as to make them insignificant. A theory unable to make sense of them is fundamentally deficient.

4 Prospects

Supposing that the redundancy argument goes through, what's next? What's an alternative to the objectivist views we have examined?

A clear beneficiary of the redundancy argument is 'prohibitionism' with its claim of the synonymy between slurs and their neutral counterparts.¹⁵ On this view, a slur is just a stylistic variation of its neutral counterpart which, for social and moral reasons, it is prohibited to use. The utterances 'A is a kraut' and 'A is a German' have exactly the same semantic meaning. The only difference between them is rather like a difference between 'B is a horse' and 'B is a steed': it is a difference in tone, a difference explained by miscellaneous experiences that the speakers and the audiences may have with these utterances.

Prohibitionism can easily deal with the High-frequency and Alternating usages. Among Germanophobes 'kraut' is often just a way of referring to Germans, and so is 'German'. On other occasions, Germanophobes may use 'kraut' just as successfully to convey various anti-German prejudices. With appropriate tweaks of the common ground, gesture, or intonation, 'German' will be used to the same end. A slur, in other words, may be an effective instrument to convey derogatory content, but only in certain contexts of use. By the same token, the natural signaling view considered earlier

¹⁵Prohibitionism: Anderson and Lepore (2013), Lepore and Stone (2018).

would be at home with prohibitionism: the use of ‘kraut’, like the use of ‘steed’, would often be a reliable cue of the speaker’s milieu.

The synonymy of slurs and their neutral counterparts does not entail that their respective uses are equally frequent or equally appropriate. This is the case with other synonyms too: reports delivered by various experts and officials would consistently differ in style from the reports by the uninitiated. Experts and officials may further insist that theirs is the ‘correct’ style, but they are unlikely to also insist that the uninitiated describe incorrectly how things are. If one tone is offensive, and another pathetic, this is a matter of psychology, not of meaning.

Prohibitionists develop this line of thought by appealing to the notion of taboo. Slurs are not just any ordinary synonyms of the neutral terms. There is an outright ban on their use. Because slurs evoke in the targeted groups (and in non-bigots generally) various negative experiences, because their use threatens public order, and for all sorts of other sundry reasons, speakers are prohibited from using them. Nothing in this ban has to do with the semantic properties of the term.

It may be thought, however, that the redundancy argument offers no succour to prohibitionism. This is because, according to some of its critics, prohibitionism is beyond salvation and should be disqualified from the start. Since, according to prohibitionism, slurs and neutral counterparts are synonymous, *a fortiori* they are co-extensional. But, the critics charge, this is false.¹⁶ Having the extension of a neutral term fixed, the extension of the relevant slur may both contract and expand. Thus the following utterances are felicitous:

- (6) a. A is a Jew, but he is not a kike. (Contracting use)
 b. B is a faggot! [spoken of a manifestly heterosexual male] (Expanding use)

According to the critics, utterances like (6) show that the extensions $\llbracket \text{KIKE} \rrbracket$ and $\llbracket \text{JEW} \rrbracket$, $\llbracket \text{FAGGOT} \rrbracket$ and $\llbracket \text{HOMOSEXUAL} \rrbracket$ don’t match. Therefore, until and unless expansion and contraction are accounted for, prohibitionism has nothing to go for it.¹⁷ Another prominent objection concerns the explanatory role of the taboo. Derogation, the critics say, must explain prohibition, not vice versa. Thus there is no escape from attributing to slurs derogatory power (i.e. associating them with a derogatory content), and the idea of taboo is explanatorily idle.¹⁸ Let me outline responses to these three charges (the details are elaborated at a greater length in the companion paper).

(I) To begin with contracting use, consider:

- (7) A: Matt is a Jew, but he is not a kike (though most other Jews *are* kikes).

The critics take an utterance like (7) as a key piece of evidence for the claim that, according to A, the extension $\llbracket \text{KIKE} \rrbracket$ is a proper subset of $\llbracket \text{JEW} \rrbracket$.

¹⁶See Croom (2015), Neufeld (2019), Falbo (2021). See Hom and May (2013) for another version of this complaint. Generally, it is directed not only at prohibitionism, but against all so-called ‘identity theories’ that treat slurs and their neutral correlates as co-extensive.

¹⁷See Neufeld (2019) and references therein.

¹⁸See Sennet and Copp (2020), Popa-Wyatt and Wyatt (2018) for this line of argument.

Well, suppose we ask the putative bigot *A* why he uttered (7). A likely answer is that he absolutely refuses to call Matt a ‘kike’ (‘Matt saved my life, how dare you!’), he says). Then this is what it is—the refusal to *call* Matt a ‘kike’:

Matt is a Jew, but he is not a ‘kike’ (though most other Jews *are* properly called ‘kikes’).

The speaker refuses to assert ‘Matt is a kike’, although he doesn’t necessarily reject the truth of that statement. Some evidence of this may be drawn from the paraphrases of (7) that are incongruous to my ear anyway:

[?] That’s a *mistake*: Matt is a Jew, not a kike.

[?] It is simply *false* that Matt is a kike. He is a Jew, though.

It seems improbable that the dispute between the speaker and the interlocutor turns on the truth of the original statement. It is more naturally construed as being about the choice of words. We observe a similar phenomenon with idioms:

(A:) — The grandpa kicked the bucket.

(B:) — The grandpa didn’t ‘kick the bucket’, he just died!

B is not saying that to die is one thing, to kick the bucket is another. Conceivably, *B* merely complains about *A*’s way of speaking.¹⁹

By the same token, the speaker in (7) complains about the style of describing Matt, not its factual accuracy. Just as the grandpa, strictly speaking, ‘kicked the bucket’, so is Matt a ‘kike’. Neither of them, however, should appropriately be described so. I conclude that the evidence for contracting use is not compelling.

(II) Turning to expanding use, compare:

- (8) a. *A*: My brother Luke is a true kike! Not because he is Jewish—he isn’t!—but because he’s so greedy!²⁰
 b. *B*: My brother Mark is a true pig! Not because he isn’t human—he is!—but because he’s so filthy!

The critics take (8a) to confirm that, according to *A*, [KIKE] contains at least one non-Jewish individual. But how different is it from (8b)? This utterance too may be taken to show that, according to *B*, [PIG] contains at least one human. But that’s not how we normally think of (8b). We think that *B* came up with a metaphor, never meaning to challenge the distinction between (animal) pigs and humans.

To spell this out a bit: Metaphorical statements aim at a novel classification of things.²¹ Before, *x* was classified as *F*. It was also routinely assumed that the classes of *F*s and *G*s don’t overlap. We now reclassify *x* overtly as *G*, whilst challenging neither the prior, familiar classification of it as *F*, nor the no-overlap assumption.

¹⁹The same applies to the use of diminutives and slang, as I explain in the companion paper.

²⁰(8a) is adapted from Croom (2015:32). The present author witnessed just this sort of use on a couple of occasions.

²¹Metaphor and classification: Goodman (1968:68ff), Black (1962:38ff), Stern (2000:153ff), Camp (2015:50–53).

Hence metaphorical predication results in a conflict of expectations and a sense of surprise. The conflict is resolved and surprise alleviated by rethinking the role of the classification of x as G . Thus, if ‘kike’ was coined to refer exclusively to Jews, applying it to a particular non-Jewish x is to classify x in terms of a salient property like greed commonly ascribed by bigots to Jews. Aided by a network of features associated with Jews in the bigoted conception of them, we are encouraged to *see* x as exemplifying some of these features salient on the occasion. Similarly, applying ‘pig’ to a human x is to think of x in terms of filthiness commonly ascribed to animal pigs. It is to see him primarily as a filthy person.

There is, therefore, good evidence that an utterance like (8a) is accounted for by metaphorical use. If the critic insists that, e.g., a non-Jewish person may felicitously be called a ‘kike’, the reply in short is that the same may be said of the metaphor in (8b). Therefore, unless the critic demonstrates the essential difference between (8a) and (8b), or unless we are prepared to ignore the metaphorical employment of terms, no proof of expanding use is forthcoming.

(III) The concept of taboo may play a vital explanatory role if, in conformity with Nunberg’s hypothesis endorsed earlier, we assume that there is a community of bigots frequently and deliberately using slurs (i.e. what *we* call ‘slurs’) among themselves. Their linguistic usage is what distinguishes them from non-bigots. Bigots, we conceive, are guilty of a moral violation. That is why non-bigots wish to segregate themselves from bigoted communities. Thus, e.g., someone who uses a racist slur among non-bigots behaves *like* a racist. In doing so, he breaks the boundary between the non-racist and the racist groups.

On this view, the special fault we find with the use of slurs is that boundary violation, and not the alleged derogation or offence. This explains why we harshly condemn their use even when the speaker does not mean to offend (as may be the case with Aldo Raine). This is not to deny that the reason why the taboo on slurs has emerged is in the behaviour and attitudes of the relevant groups of bigots, in what may broadly be described as their derogation of the targeted groups. But, having separated themselves from bigots, the non-bigots insist on the inviolable boundary between them. They want to ‘have nothing to do’ with bigots. A further related source of our reactions to the use of slurs is in ‘magical contagion’ recently explored in the literature of disgust (even if we can’t reconcile our reactions with our prevailing modern outlook). Bigotry, like many grave moral faults, is seen as contagious. Thus the slur user puts the audience of non-bigots in danger of becoming bigots themselves. His interlocutors recoil from the use of slurs, since they are ‘contaminated’ by it.

This explanation accords well with the classical accounts of taboo in early societies. By taking the idiom of ‘taboo’ more seriously than the prohibitionists have done so far we can also explain better our reactions to the use of slurs in embedded constructions. The explanatory job is done by the same notion of contagious contamination. Even if a speaker does not himself call anyone a ‘kike’ or a ‘faggot’, we feel that we are contaminated with bigotry when we are exposed to the use of these terms in indirect speech or under negation. In the companion paper I substantiate this claim by linking it to the psychological theories of disgust and ‘magical thinking’ (I also address there the special case of direct quotation).

5 Review

To sum up: the superficially plausible idea that slurs are associated with derogation leads to the violation of the maxim of quantity (the redundancy of derogation). This is so on the assumption (‘Nunberg’s hypothesis’) that slurs are routinely used in the in-group setting. While violations of the maxim of quantity may be justified in some types of conversations, I have argued that none of them fits the case of slurs. Finally, I briefly argued that, if the redundancy argument goes through, prohibitionism should benefit from it at the expense of the ‘objectivist’ views.

References

- Anderson, L., & Lepore, E. (2013). Slurring words. *Noûs*, 47(1), 25–48.
- Arkes, H. P. (1974). Civility and the restriction of speech. *Supreme Court Review*, 281–336.
- Black, M. (1962). *Models and metaphors*. Ithaca: Cornell University Press.
- Bolinger, R. J. (2017). The pragmatics of slurs. *Noûs*, 51(3), 439–462.
- Brown, P. (2000). Repetition. *Journal of Linguistic Anthropology*, 9(1–2), 223–226.
- Camp, E. (2013). Slurring perspectives. *Analytic Philosophy*, 54(3), 330–349.
- Camp, E. (2015). Why metaphors make good insults: perspectives, presupposition, and pragmatics. *Philosophical Studies*, 174(1), 47–64.
- Camp, E. (2018). A dual act analysis of slurs. In D. Sosa (Ed.) *Bad words*. New York: Oxford University Press.
- Croom, A. M. (2015). The semantics of slurs: a refutation of coreferentialism. *Ampersand*, 2(1), 30–38.
- Davies, C., & Arnold, J. E. (2019). Reference and informativeness. In C. Cummins, & N. Katsos (Eds.) *The Oxford handbook of experimental semantics and pragmatics*. Oxford University Press.
- Engelhardt, P. E., Bailey, K. G. D., & Ferreira, F. (2006). Do speakers and listeners observe the Gricean maxim of quantity? *Journal of Memory and Language*, 54(4), 554–573.
- Falbo, A. (2021). Slurs, neutral counterparts, and what you could have said. *Analytic Philosophy*, 00, 1–17. Forthcoming.
- Fox, D. (2014). Cancelling the maxim of quantity. *Semantics and Pragmatics*, 7(5), 1–20.
- Goodman, H. N. (1968). *Languages of art*. Indianapolis and New York: The Bobbs-Merrill Company.
- Grice, H. P. (1989). Logic and conversation. In *Studies in the way of words*. Harvard University Press.
- Hom, C. (2008). The semantics of racial epithets. *Journal of Philosophy*, 105(8), 416–440.
- Hom, C. (2010). Pejoratives. *Philosophy Compass*, 5(2), 164–185.
- Hom, C. (2012). A puzzle about pejoratives. *Philosophical Studies*, 159(3), 383–405.
- Hom, C., & May, R. (2013). Moral and semantic innocence. *Analytic Philosophy*, 54(3), 293–313.
- Horn, L. R. (2004). Implicature. In L. R. Horn, & G. Ward (Eds.) *The handbook of pragmatics* (pp. 3–28). Oxford: Blackwell.
- Hornsby, J. (2002). Meaning and uselessness. In P. French, & H. Wettstein (Eds.) *Midwest studies in philosophy*, (Vol. 25 pp. 128–141). Oxford: Blackwell.
- Jeshion, R. B. (2018). Slurs, dehumanization, and the expression of contempt. In D. Sosa (Ed.) *Bad words* (pp. 77–107). New York: Oxford University Press.
- Kennedy, R. L. (2002). *Nigger*. Vintage Books.
- Lepore, E., & Stone, M. (2018). Pejorative tone. In D. Sosa (Ed.) *Bad words* (pp. 132–154). New York: Oxford University Press.
- Liu, C. (2021). Slurs as illocutionary force indicators. *Philosophia*. Forthcoming.
- Neufeld, E. (2019). An essentialist theory of the meaning of slurs. *Philosophers’ Imprint*, 19, 35.
- Nunberg, G. (2018). The social life of slurs. In D. Fogal, D. W. Harris, & M. Moss (Eds.) *New work on speech acts* (pp. 237–295). New York: Oxford University Press.
- Popa-Wyatt, M., & Wyatt, J. L. (2018). Slurs, roles and power. *Philosophical Studies*, 175, 2879–2906.
- Potts, C. (2015). Presupposition and implicature. In S. Lappin, & C. Fox (Eds.) *The handbook of contemporary semantic theory* (pp. 168–202). Oxford: Wiley Blackwell.
- Rubio-Fernandez, P. (2019). Overinformative speakers are cooperative. *Cognitive Science*, 43(11), e12797.

- Schegloff, E. A. (1996). Confirming allusions. *American Journal of Sociology*, *102*(1), 161–216.
- Sennet, A., & Copp, D. (2020). Pejorative verbs and the prospects for a unified theory of slurs. *Analytic Philosophy*, *61*(2), 130–151.
- Skyrms, B. (2010). *Signals*. Oxford University Press.
- Stern, J. J. (2000). *Metaphor in Context*. Cambridge: The MIT Press.
- Tannen, D. (2007). *Talking voices*, 2nd edn. Cambridge University Press.
- Torrenzo, G. (2020). Slurs and semantic indeterminacy. *Philosophia*, *48*, 1617–1627.
- Whiting, D. (2013). It's not what you said, it's the way you said it: Slurs and conventional implicatures. *Analytic Philosophy*, *54*(3), 364–377.
- Williamson, T. (2009). Reference, inference, and the semantics of pejoratives. In J. Almog, & P. Leonardi (Eds.) *The philosophy of David Kaplan* (pp. 137–158). New York: Oxford University Press.

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