



Generic View of Gendered Slurs and the Subset Argument

ABSTRACT: *The neutral counterpart assumption is widely accepted in the study of slurs. It provides a simple and effective explanation for the meaning of slurs. Slurring terms are coextensional with their neutral counterparts. However, Lauren Ashwell (2016) has questioned this assumption. She argues that gendered slurs refer to a subset of their neutral counterparts. Hence, slurs are not coextensional with their counterparts. She goes on to present a view that is not based on the counterpart assumption. Still, her view is a unifying view of slurs as it also applies to ethnic and racial slurs. In this paper, I defend the counterpart assumption with a generic view of slurs. While being a unifying view, it accommodates the subset argument with its eponymous feature that the meaning of slurs involves a generic component.*

KEYWORDS: gendered slurs, ethnic and racial slurs, neutral counterpart assumption, generics

I. Counterpart Assumption and the Subset Argument

I.1 Counterpart Assumption

The idea that slurs have neutral counterparts is widely accepted. Let us call it the *neutral counterpart assumption* or the *counterpart assumption* for short. It provides a very simple and effective explanation for the meaning of slurs. A neutral counterpart of a slur is a nonderogatory coextensional expression. For example, the semantics of ‘Boche’ is parasitical to the semantics of its neutral counterpart. The truth conditional contribution of ‘Boche’ is identical to the truth conditional contribution of ‘German’. It also provides an application criterion for slurs: ‘Boche’ is applied correctly to the set of German people. (In section 2.1, the thesis concerning semantics and the thesis concerning application are separated. Together they form a stronger version of the counterpart assumption while the latter forms a weaker version of the counterpart assumption. But at this stage, I will go with the stronger version because Lauren Ashwell [2016] targets that thesis.)

The counterpart assumption yields an equally powerful explanation for the derogation associated with slurs. Slurs are especially derogatory because they

I thank everyone in philosophy at Tampere University for a warm and friendly atmosphere. I also thank Giulio Pietrousti. The idea for the expressive linguistic test stemmed from his comments while I was a visiting member of LOGOS research group. Many thanks to the referees for helpful comments and suggestions. This work is supported by the Kone Foundation. I am also grateful for the opportunity to work in Kristiina Rolin’s project ‘Social and Cognitive Diversity in Science’ funded by the Academy of Finland.

derogate the target solely on the basis of group membership. This membership is specified with the neutral counterpart. Personal merits or demerits are irrelevant. Slurs may denigrate individuals, for example, on the basis of the color of their skin or some other feature that associates them with demographic groups.

1.2 Subset Argument

In her ‘Gendered Slurs’ (2016: 228–39), Lauren Ashwell questions the counterpart assumption with what I call a *subset argument*. She argues that gendered slurs do not have neutral counterparts. Ashwell quite rightly points out that gendered slurs are mentioned only in passing in the literature and it is generally assumed that the neutral counterpart of ‘slut’ is ‘woman’. However, according to Ashwell, while ‘slut’ does indeed refer to women, it does not refer to all women. Rather, it refers to women who are promiscuous. As a consequence, ‘slut’ refers to a proper subset of women. The natural next question is whether the description ‘woman who is promiscuous’ could be the counterpart. Perhaps gendered slurs do not have common nouns like ‘German’ as counterparts but rather descriptions. Nevertheless, Ashwell rejects the idea of a descriptive counterpart:

‘woman who has sex with a lot of sexual partners’ is in fact not co-extensive with ‘slut.’ Rather than just being about a sheer number of partners, ‘slut’ is often applied to someone because of who those partners are and what she does with them. (2016: 235)

She continues:

Instead, the closest phrase to a true correlate for ‘slut’—albeit one that is certainly not neutral—is something that attributes a disposition toward sexual behavior that is deemed inappropriate; perhaps: ‘woman who is inappropriately disposed toward sexual relations,’ [. . .] (2016: 235)

As a result, Ashwell proposes a dispositional view of the content of ‘slut’; *mutatis mutandis* for other gendered slurs. For example, the meaning of ‘bitch’ is something like ‘disposed to be more boisterous, more assertive, more self-concerned than is appropriate for a woman’ (Ashwell 2016: 235). From these formulations, we can see the challenge. The commonality with gendered slurs is that the additional description that narrows the extension is not purely descriptive. It is also evaluative and normative. In comparison to ethnic and racial slurs, gendered slurs do not have *neutral* counterparts like African American or Jewish.

1.3 Unifying Normative View of Slurs

From this observation, Ashwell goes on to propose a unifying view of slurs, according to which the associated derogation stems from their normative aspects and not from the counterpart relation even if a given slur might have a neutral

counterpart. She argues that derogation is based on *normative constraints* on the target group. Slurs set up questionable norms on how people ought to act (see Diaz Legaspe 2018: 243). This applies to gendered slurs as ‘slut’ and ‘bitch’ demonstrate. In addition, the normative constraint also applies to ethnic and racial slurs. She refers to Chris Rock’s stand-up line:

(1) I love black people, but I hate niggers.

Ashwell argues that (1) makes sense only if the N-word refers to a proper subset of black people, to those who act or are disposed to act in a way that is inappropriate according to racist standards. This explanation for derogation does not appeal to the counterpart relation and therefore ‘we ought to rethink whether this is essential to any slur’. Derogation can be explained without the assumption. Furthermore, the criterion for the application can be set up without the counterpart assumption: if the N-word sets up racist norms for how black people ought to act and picks out those who fail to meet these norms, then applying the N-word to a white person is a linguistic mistake (Ashwell 2016: 239).

In sum, Ashwell first puts forward the subset argument that questions the counterpart assumption. After this, she formulates a unifying theory of slurs that does not rely on the counterpart assumption. Instead, her view relies on the normative constraints on the target group. In the following, I will first present Justina Diaz Legaspe’s thoughts on the subset argument. Her response aims to accommodate the subset argument while retaining the counterpart assumption. However, she does not develop a unifying view. In fact, she argues that the semantic account of gendered slurs is different from racial and ethnic slurs. I will then take on board some of Diaz Legaspe’s insights as I develop my own response to the subset argument. My generic view is a unifying view. It defends the idea that slurs, whether ethnic and racial or gendered slurs, are coextensional with their counterparts. I argue that the evidence supporting the subset argument does not lead to the rejection of the counterpart assumption. Rather, the evidence suggests that the meaning of all slurs involves a generic element.

2. Diaz Legaspe and the Subset Argument

2.1 Specified Counterpart Assumption

In her ‘Normalizing slurs and out-group slurs’ (2018: 234–55), Diaz Legaspe discusses Ashwell’s argument at length. She accepts that gendered slurs do indeed refer to a proper subset of the target set. From this perspective, she develops a view of slurs that can accommodate the referential restriction while adhering to the counterpart assumption. Her key insights clarify the counterpart assumption decisively.

Diaz Legaspe is not entirely convinced by Ashwell’s explanation for derogation. She points out that Ashwell’s view fails to account for the full derogatory force of slurs. Even though yelling ‘You are not acting according to what is expected from black folks!’ will most likely generate a heated discussion about how exactly black

people should behave, it is certainly not as derogatory as yelling the N-word. Diaz Legaspe thinks that the counterpart assumption is still a far more promising explanation for derogation (2018: 243–50). First, she specifies the counterpart assumption in an illuminating way. The counterpart assumption breaks down to two different versions, the *Application Neutral Counterpart Thesis*:

(Application Thesis) For every slurring expression e there is a neutral counterpart NC_e and the correct application criteria for NC_e are identical to the correct application criteria for e .

and the *Referential Neutral Counterpart Thesis*:

(Reference Thesis) For every slurring expression e there is a neutral counterpart NC_e and the class of individuals referred to by NC_e (call it $[[NC_e]]$) is identical to the class of individuals referred to by e .

Application Thesis states that the application of a slurring expression e is correct only within the set determined by the neutral counterpart NC_e . The thesis then answers the question of ‘to which people the slurring term e applies’ (Diaz Legaspe 2018: 234–36). However, it seems unlikely that Application Thesis provides the criteria for full competency involving the slurring term e . As Williamson says, a fully competent speaker must understand both the truth conditional contribution and the derogatory aspect of e (see Williamson 2009: 152–53). If the derogatory aspect is included in the criteria concerning application, then the application conditions for the slurring term e cannot be identical with NC_e . Hence, it has to be emphasized that Application Thesis yields only an answer to the above question. In any case, Application Thesis is the weaker of the two because it is neutral regarding the truth conditional contribution of slurs. In contrast, Reference Thesis defines the truth conditional contribution of slurs *via* the neutral counterparts (Diaz Legaspe 2018: 234–36). For example, we can use Reference Thesis to define the extension of ‘Boche’. Reference Thesis states that the extension of ‘Boche’ is identical to the semantic value of its neutral counterpart NC_{Boche} . We then have

$$[[NC_{Boche}]] = \{x | x \text{ is German}\}$$

that states that the semantic value of the neutral counterpart of ‘Boche’ is the set of German people. Hence, the extension of ‘Boche’ is the set of German people.

Reference Thesis also entails Application Thesis. As far as I know, anyone who adheres to Reference Thesis is also committed to Application Thesis, but not the other way round. For example, Christopher Hom and Robert May adhere to Application Thesis but they reject Reference Thesis. They agree that ‘Boche’ applies only to the German people but reject the idea that the truth conditional contribution of ‘Boche’ and ‘German’ is the same. Hom and May think that while ‘German’ refers to the set of German people, the reference of ‘Boche’ is an empty set. Nevertheless, they do think that Application Thesis is related to semantic

content of slurs. For example, the content of ‘Boche’ is something like ‘ x ought to be the target of negative moral evaluation because of being German’ but no one deserves a negative moral evaluation because of their racial or ethnic background. Even though the reference is an empty set, the content of ‘Boche’ still involves the Germans (Hom and May 2013: 294–300.)

It is also important to point out that Ashwell’s target is clearly Reference Thesis. Specifically, her target is the idea that Application thesis is entailed by Reference Thesis. As seen in the earlier quotation, she thinks that because ‘slut’ does not have a neutral coextensional counterpart, the counterpart assumption cannot be right. In fact, Ashwell’s view is compatible with Application Thesis. In her view, the N-word sets up norms for black people and so to apply it to a white person would constitute a linguistic mistake (Ashwell 2016: 235–39). Hence, the application can be explained without Reference Thesis.

Diaz Legaspe goes on to make her own proposal based on a distinction between the semantics of ethnic and racial slurs and the semantics of gendered slurs. Reference Thesis applies to ethnic and racial slurs. This means that, strictly speaking, (1) is an infelicitous sentence because Rock clearly intends to target a subset of black people. Diaz Legaspe explains that sometimes language can be used in an idiosyncratic way and that this is one of those instances. In Rock’s own idiolect, the N-word refers to a subset of black people, and the rest of us can understand Rock’s idiolect. However, this does not change the semantics of ethnic and racial slurs in our common language; Reference Thesis still explains their semantics. In contrast, given the linguistic evidence concerning the use of gendered slurs, Diaz Legaspe admits that Reference Thesis does not explain the semantics of gendered slurs because they refer to a subset of the target gender: ‘any competent speaker of English would agree with the fact that “slut” is used to refer to women that behave in a certain way or that seem prone to behave in a certain way’. There is, then, a distinction between the *uses* of ethnic and racial slurs with restricted reference and the *meanings* of gendered slurs based on restricted reference. Ethnic and racial slurs can *sometimes* be used with a referential restriction, as (1) illustrates, but the reference of gendered slurs is *always* a proper subset of the target group (Diaz Legaspe 2018: 344–46).

The target of Ashwell’s argument is Reference Thesis. Diaz Legaspe’s strategy, therefore, is to maintain the distinction between ethnic and racial slurs and gendered slurs and, at the same time, to come up with a modified referential thesis for gendered slurs that is still compatible with the counterpart assumption. First, the narrowing of the extension is down to stereotypes. They determine the extension to be a subset of the target gender. In this context, Diaz Legaspe introduces *P-behavior*: gendered slurs ‘target the members of that class that exhibit a certain property: a certain behaviour or disposition to behave in a certain way’ (Diaz Legaspe 2018: 246). (In section 4.2, the role of stereotypes in derogation is explored in detail.) To cope with this semantic fact, Diaz Legaspe modifies Reference Thesis. While the original Reference Thesis claimed that the semantic value of a slurring term e is identical to the semantic value of NC_e , this does not hold for gendered slurs. Instead, a *Restricted Reference Thesis* does hold:

(Restricted Reference) Whenever ‘ o is an e ’ is true, ‘ o is a NC_e ’ is also true.

On the contrary, the converse of Restricted Reference does not hold. You cannot say that whenever o is a NC_e , o is also an e because the semantic value of a slurring term e is a proper subset of the semantic value of NC_e . With ethnic and racial slurs, these kinds of two-way inference patterns hold because the semantic values are identical.

According to Restricted Reference, the neutral counterpart still has a role in determining the semantic value of the corresponding slurring term. Diaz Legaspe concludes that the above observations support the idea that gendered slurs have an associated neutral counterpart, even though slurs and their counterparts are not coextensional (Diaz Legaspe 2018: 249).

2.2 Assessment of Diaz Legaspe’s Proposal

Diaz Legaspe argues that even though gendered slurs and their counterparts are not coextensional, Restricted Reference shows that NC_e still has a role in determining the reference of the gendered slurring term e . Furthermore, NC_e provides application conditions for e . Only women can be called ‘sluts’.

Although Diaz Legaspe provides valuable insight into the nature of the counterpart assumption, I have two concerns. First, you might ask how exactly her proposal differs from Ashwell’s view. After all, Ashwell’s view is compatible with Application Thesis and her rejection of the counterpart assumption rests on the fact that the semantics of gendered slurs are different from their counterparts. Diaz Legaspe’s proposal ends up saying the same thing: while the counterpart relation does provide an application condition for gendered slurs, her proposal still admits that the semantics of gendered slurs differ from the semantics of their counterparts. Second, Ashwell provides a unifying view of slurs but Diaz Legaspe does not. According to Ashwell, derogation is not based on the counterpart assumption but on the normative constraint. Even though ‘slut’ refers only to a subset of women, the normative constraint applies to all women. Furthermore, the normative constraint is the commonality in slurs across the board. Diaz Legaspe’s *contrasting view* holds that there is a difference between gendered slurs and ethnic and racial slurs. The difference is in the meaning of slurs: the semantic value of ethnic and racial slurs is identical to that of their neutral counterparts whereas the semantic value of gendered slurs differs from that of their counterparts. (In section 4.3, the difference between unifying and contrasting views is examined in detail.)

3. Generic View of Slurs

3.1 Genericity of Ethnic and Racial Slurs

My proposal aims to defend the counterpart assumption against the subset argument by providing a unifying semantics of slurs that rests on the counterpart assumption. The subset argument is supported by the fact that there is an extra feature involved

with gendered slurs, the P-behavior, that then narrows the reference to a proper subset of the target set. According to my generic view, this extra feature is not unique to gendered slurs. Ethnic and racial slurs also add a negative stereotype to the target. This stereotype is attributed to the target in a generic way. As a result of this genericity, the application of slurs allows counterexamples, even with ethnic and racial slurs. This accommodates the subset argument. At the same time, the counterpart assumption holds.

The proposal builds on a view of ethnic and racial slurs introduced and detailed in my ‘Generic inferential rules for slurs: Dummett and Williamson on ethnic pejoratives’ (Valtonen 2021: 198, 6533–51). Here the view is expanded to cover gendered slurs. Before applying the view to gendered slurs, I will briefly go through the relevant features of the view. The proposed view is based on an inferentialist framework originally introduced by Michael Dummett. According to Dummett, slurs are not only morally objectionable, but they are also proof-theoretically objectionable. According to inferentialism, the inferential rules for terms determine the semantics for language. Conversely, the referential direction starts with the semantics and determines the valid inferential rules on the basis of that semantics. Dummett contends that the condition for the application of ‘Boche’ is that the target is German but the consequence of the application is that the target is ‘barbarious and more prone to cruelty than other Europeans’ (Dummett 1973: 454–55). Let us formulate the rules for ‘Boche’ as

$$\text{Boche-I: } \frac{x \text{ is German}}{x \text{ is Boche}} \quad \text{Boche-E: } \frac{x \text{ is Boche}}{x \text{ is cruel}}$$

Here, the introduction rule (Boche-I) determines the conditions for the application of ‘Boche’ and the elimination rule (Boche-E) determines the consequences of the application. The rules are proof-theoretically bad because they are nonharmonious. Boche-E unpacks more than Boche-I packs in, as Ian Rumfitt puts it (Rumfitt 2000: 789).¹ The rules are also epistemically bad because they allow the attribution of cruelty to the German people without evidence. If inferentialism is right, then the inferential rules explicate the meaning of ‘Boche’. As a result, it is the meaning of ‘Boche’ that licenses the attribution of cruelty, not the evidence. This then makes slurs epistemically objectionable. However, in his critical assessment, Timothy Williamson points out that the rules are not compatible with the counterpart assumption. The counterpart assumption says that ‘Boche’ applies only to German people and ‘slut’ only to women. The problem with Boche-I and Boche-E is that Boche-I assigns the reference to be the set of German people, while Boche-E assigns the reference to be the set of cruel people. Hence, the assignment of reference for ‘Boche’ is the union of German people and cruel people but this is not compatible with the counterpart assumption. Although Stalin was cruel, he was not a ‘Boche’. He was Russian (Williamson 2009: 142–48).

¹ In comparison, the rules for conjunction are harmonious: &-I: A, B ⊢ A & B; &-E(1): A & B ⊢ A; &-E(2): A & B ⊢ B. The elimination rules get you back where you started with the introduction rule.

In the following, the aim is to come up with rules for ‘Boche’ that are compatible both with Application Thesis and Reference Thesis. The proposed inferential rules aim to make ‘German’ and ‘Boche’ coextensional. They also honor the Dummettian insight that slurs are epistemically objectionable. In turn, this will lead to a unifying view of slurs that can accommodate the supporting evidence for the subset argument. I will first discuss the view in relation to ethnic and racial slurs and then move on to gendered slurs.

Williamson remarks that the iniquity in ‘Boche’ is most likely in a generic form, something like ‘There is a tendency for Germans to be cruel’ (Williamson 2009: 151). Dummett’s original rules may also include a generic element. His original rules allowed the conclusion that Germans are prone to cruelty, which seems like a generic formulation. In any case, what needs to be explained is how the generic element helps to overcome the incompatibility between the counterpart assumption and the inferential rules. I propose the following generic rules for ‘Boche’:

$$\text{Boche-I(Gen): } \frac{x \text{ is German}}{x \text{ is a Boche}}$$

$$\text{Boche-E(Gen): } \frac{x \text{ is a Boche}}{x \text{ is German and Germans are typically cruel}}$$

More formally, Boche-E can be formulated as follows:

$$\text{Boche-E(Gen): } \frac{x \text{ is a Boche}}{x \text{ is German \& Gen } x \text{ [German}(x)\text{][cruel}(x)}}$$

This formalization is a pretty standard way to represent the structure of generic sentences. ‘Gen’ is a generic operator comparable to quantifiers, but unlike quantifiers, it does not specify the exact relationship between the sets in the scope of the operator. In natural language, the generic operator can be substituted with terms like ‘usually’, ‘generally’, or ‘typically’; although it is often omitted altogether (see Leslie 2008: 1–6).

There are two distinguishing features to generics: we assent to them and the assent holds in the face of counterexamples. The curious thing about generics is that the proportion of individuals needed to confirm a generic varies wildly. Take the generic ‘Tigers are striped’. Here the proportion of relevant individuals is very high. Only very few tigers are stripeless. With ‘Ducks lay eggs’ the share of relevant individuals is around 50% since only female ducks lay eggs. This is still relatively high compared to ‘Mosquitoes spread West Nile virus’. In actuality, less than 1% of mosquitoes spread the virus. Yet, we tend to assent to the sentence. The study of generics usually takes a semantic approach that seeks to identify truth conditions for ‘Tigers are striped’, ‘Ducks lay eggs’ and ‘Mosquitoes spread West Nile virus’. These truth conditions would then confirm the sentences to be compositionally true despite the variation in the number of relevant individuals.

This points to an obvious difficulty concerning Boche-I(Gen) and Boche-E(Gen). ‘Germans are typically cruel’ is not true. So a semantic view will not pass muster because it cannot accommodate the counterpart assumption. According to Williamson, this can be achieved only by assigning the set of Germans as the reference of ‘Boche’ but the semantic interpretation of generics results in indeterminacy of reference. Boche-I(Gen) fixes the reference to the set of Germans but Boche-E(Gen) determines the reference to an empty set since the conjunction in Boche-E(Gen) is false. This is a rather odd result. On the one hand, ‘Boche’ refers to German people. On the other hand, it refers to an empty set. This is hardly compatible with the counterpart assumption. Luckily, Sarah-Jane Leslie proposes an alternative approach to generics. She is sceptical whether any semantic view can explain the puzzle of generics. Leslie argues that these kinds of ‘generalizations do not operate on set extensions’. ‘They are not grounded in such extensional or statistical information [. . .]’, she adds (2007: 394). She proposes a psychological view instead, hypothesizing that generics can be explained with a psychological mechanism of generalization. The mechanism is by design an efficient information-gathering mechanism and a very basic way of obtaining information about environment. (2007: 383–84; 2008: 18–23.)

Semantic views hold that the key to generics is to give an analysis for the Gen-operator. This would then yield semantic truth conditions for individual generics. In response, Leslie first admits that ‘Tigers are striped’ should be represented as ‘Gen x [Tiger(x)] [Striped(x)]’ but then refrains from any further analysis of generics. In her view, there are no *semantic* truth conditions for generics in the sense that the Gen-operator would contribute compositionally to the truth conditions of ‘Tigers are striped’. Instead, she offers much looser ‘worldly truth-makers’ for generics. She asserts that the mechanism of generalization involves three types of generalizations, each tracking a different type of truth-maker in the world. First, there is a characteristic dimension that aims to identify regularities in the world. It seeks characteristic features of kinds. Animals of the same kind tend to make similar noises, for example. Hence, ‘Ducks lay eggs’ is true because characteristically ducks lay eggs and the counterexamples are only negative. That is, the male ducks do not exhibit alternative ways of reproduction. Second, majority generics are true simply iff the majority possesses the attributed feature. For example, ‘Tigers have stripes’ is true because the majority of tigers do have stripes. Finally, the mechanism registers and generalizes information that is striking. Leslie claims that the striking feature is often horrific or appalling. Hence, generics like ‘Mosquitoes spread West Nile disease’ and ‘pitbulls maul children’ are considered true because both attribute a pretty horrific property to the subjects even though only a tiny proportion of mosquitoes actually spread the disease and pitbulls rarely maul anyone. Nonetheless, the striking feature generalization often acts as a good predictor of a property that may pose a threat to the agent. Leslie also mentions that the erroneous generalization ‘Muslims are terrorists’ is most likely a product of the striking feature generalization (2007: 383–86). This remark more or less settles the question of what type of generalization ‘Germans are cruel’ is. Just as only a small proportion of mosquitoes carry the West Nile virus, the virus-free

mosquitoes are still disposed to carry the virus, the bigots might think that even though not all Germans exhibit signs of cruelty, they are still disposed to act in a cruel way (see Leslie 2007: 385). Leslie says that the psychological mechanism of generalization is closely connected with psychological essentialism according to which people tend to believe that things have an internal essence that makes them the things they are (Medin and Ortony 1989: 183). Leslie has argued that psychological essentialism can have pernicious effects concerning social kinds since it can lead to generalizations like ‘Germans are cruel’ and ‘Muslims are terrorists’ (Leslie 2017: 393–421). A Swedish murderer is perceived just as a bad apple whereas a murdering German manifests cruelty inherent to all Germans (see Leslie 2007: 385). Needless to say that the worldly truth-makers do not support this generalization.

The current construal of generics aims to codify the inferential rules for ‘Boche’ that assign the set of Germans as the reference since the objective is to defend the counterpart assumption in its strongest form. I think the developed generic rules for ‘Boche’ achieve this. Together, Boche-I(Gen) and Boche-E(Gen) determine the extension of ‘Boche’ to be the set of Germans because the generic element does not have any semantic consequences. According to Leslie’s psychological view, the generic element in Boche-E(Gen) ‘Germans are typically cruel’ does not contribute to assigning the reference to ‘Boche’. The extension is assigned with the first part of the conjunction alone (‘x is German’). The latter part of the conjunction does not contribute to this task because Gen-operator does not have an extensional interpretation. To contribute to the assignment of reference, Gen-operator must produce a clear answer to the question of ‘Over which objects Gen-operator ranges’. The semantic views aim to answer this question but, according to Leslie, Gen-operator does not have that kind of role in generic statements. It does not yield a precise set of individuals that then could be used to assign the extension of ‘Boche’. So Boche-E(Gen) assigns the reference on the basis of the first part of conjunction to the set of German people. Still, the rules honor the Dummettian insight that slurs are epistemically objectionable. Even though the latter part of the conjunction turns out to be trivial concerning the assignment of reference, it is by no means trivial concerning the inferential role of ‘Boche’. The rules confirm the idea that the conditions for the application of ‘Boche’ is that the target is German but the consequence of the application includes the attribution of cruelty to German people. It is the meaning of ‘Boche’ that licenses the claim that the Germans are typically cruel, not the evidence for the cruelty of Germans.

3.2 Derogation and Epistemic Objectionability

The generic view advances to accommodate the evidence supporting the subset argument while maintaining the counterpart assumption. There are two crucial points concerning this task. First, the bigots are very likely to say things like:

- (2) Hans is German but he’s not cruel like those Boches. He is one of the good ones.

The subset argument is based on this sort of linguistic evidence but the utterance in (2) contains an ethnic slur. Ashwell admits that, unlike gendered slurs, ethnic and racial slurs have counterparts like African American or Jewish. Diaz Legaspe believes that just because ethnic and racial slurs do have coextensional counterparts, (2) has to be idiolectic. I argue that utterances like (2) are far too common to be idiolectic. Rather, the semantics of slurs should explain them just like the generic rules do. The application of ‘Boche’ allows exceptions because the attribution of cruelty is generic in Boche-E(Gen). The second point emphasizes the connection between the current generic view and the counterpart assumption. Williamson argues in his ‘Blind Reasoning’ that the referent of ‘Boche’ cannot be a subset of Germans—that is, those Germans who actually are cruel (Boghossian and Williamson 2003: 260–61). I agree because this treatment of slurs cannot explain what is objectionable with slurs. Consider the following example:

(3) Himmler was a Boche.

According to the subset treatment, this is a true and an appropriate statement. Himmler was indeed German and cruel. However, slurs do not derogate individuals but social groups. They derogate individuals only as members of social groups. This is what is offensive and objectionable about slurs. Personal merits or demerits are irrelevant. This observation supports the idea that the reference of ‘Boche’ is the set of German people and only the Germans. That is, it highlights the connection between derogation and the counterpart assumption. It also highlights the fact that if you reject the counterpart assumption, as Ashwell does, you have to look for a new source for derogation. According to Ashwell, derogation stems from the normative constraints on the target group. I argue that the counterpart assumption coupled with the idea of epistemic objectionability provides a better overall explanation for derogation. In fact, there is a division of labor concerning the generic component and epistemic objectionability. The generic component *accommodates* the evidence put forward by Ashwell. While epistemic objectionability *emphasizes* the role of the counterpart assumption in derogation.

As already pointed out, my generic treatment of slurs takes its cue from Williamson. He says that the conventional implicature triggered by ‘Boche’ should be read as a generic ‘there is a tendency for Germans to be cruel’. According to Williamson, the nature of derogation is pragmatic and nontruth conditional. It is pragmatic because the derogatory aspect is explicated with Gricean implicature. Furthermore, the implicature does not affect the truth conditions of (3). (3) is true iff Himmler was German. According to my generic view, the nature of derogation is nontruth conditional but not pragmatic. Rather, derogation relates to Fregean senses, which are part of the truth conditional machinery. This is key in viewing slurs as epistemically objectionable. According to the Dummettian view, the inferential rules explicate the senses of the terms. Frege insisted that the senses determine the extension of terms. In short, the inferential rules determine the extension. In this sense, the inferential rules are part of the truth conditional machinery. However, since German and ‘Boche’ are coextensional, you cannot see

derogation in the truth conditions of 'Boche'. The claim that derogation is in the inferential rules allows to maintain the Dummettian insight that there is something wrong with slurs. The inferential rules for 'Boche' are objectionable because they allow to attribute cruelty to Germans without any actual evidence. They bypass the usual demands of backing your claims with evidence. To emphasize once more, it is the meaning of 'Boche' that allows cruelty, not the evidence. In contrast, the conventional implicature strategy does not say that the meaning of slurs is objectionable. Rather, it is the racist attitudes expressed with slurs that are objectionable. In my view, the epistemic objectionability highlights the contribution of the counterpart assumption in derogation. The current view explains why slurs are so offensive and demoralizing. Despite your personal moral integrity and accomplishments, with slurs, you are reduced to a criminal, to being lazy, or to a vulgarian on the basis of your skin color, religion, or some other marker of a demographic group. The generic view claims that this is done in an epistemically objectionable way. It is as if the bigots found or created a loophole in meaning enabling them to put forward objectionable ideas without the need to justify those ideas with evidence. At the same time, that loophole makes the meaning of slurs objectionable.

3.3 Genericity of Gendered Slurs

As we move on to apply this framework to gendered slurs, let us first formulate generic rules for 'slut' as:

$$\text{Slut-I(Gen): } \frac{x \text{ is a woman}}{x \text{ is a slut}}$$

$$\text{Slut-E(Gen): } \frac{x \text{ is a slut}}{x \text{ is a woman \& Gen } x \text{ [woman}(x)\text{] [promiscuous}(x)\text{]}}$$

These rules show that gendered slurs are not that different from ethnic and racial slurs. At the same time, the proposed rules are compatible with the evidence that supports the subset argument. Ashwell's thought is that 'slut' refers not only to women who behave promiscuously, but also to women who are disposed to promiscuity. The current generic rules capture this with one notable difference. According to Ashwell, the misogynist belief divides women into three categories: those who behave promiscuously, those who are disposed to promiscuity, and those who neither behave promiscuously nor are disposed to promiscuity. Ashwell thinks that 'slut' refers only to the first two categories but not to the third one. As a result, the subset argument holds (2016: 235). In contrast, in my generic view, the derogation associated with 'Boche' is based on the xenophobic idea that all Germans are disposed to cruelty even though only some Germans actively display cruelty. Similarly, 'slut' refers to women because all women are disposed to promiscuity, according to misogynist thinking. This slight difference allows the generic view to maintain the counterpart assumption. Both Application Thesis

and Reference Thesis hold as the current view adheres to Reference Thesis that then entails Application Thesis.

4. Support for the Generic View

4.1 Entitlement to Protest

Diaz Legaspe summarizes the puzzle about entitlement to protest as follows:

Now, the referential restriction of gendered slurs leaves us wondering why members of the targeted gender that do not P-behave should protest the use of the word: if ‘slut’ only refers to women who P-behave, women who do not should not be affected by its use. But in real life even women with the most impeccable behaviour should protest when some other woman is called a ‘slut’. (Diaz Legaspe 2018: 249)

My response to this is simple: gendered slurs are not constrained by the referential restriction. The counterpart assumption also applies to gendered slurs. Semantically speaking, ‘slut’ and ‘bitch’ are words that refer solely to women. Hence, the promiscuity that comes with ‘slut’ is attributed to women, albeit generically. That is why every woman is entitled to protest the use of ‘slut’. This has at least two advantages, the first over Diaz Legaspe’s proposal and the second over Ashwell’s.

As seen above, Ashwell agrees with Application Thesis and says that all women are thereby entitled to protest slurs that target women. According to her, the application criterion does not stem from the semantics of slurs. Rather, Application Thesis stems from the normative constraints. Surprisingly, Diaz Legaspe also appeals to the normative constraint. She explains that every woman is entitled to protest the use of ‘slut’ because the normative system is enforced on all women. All women are forbidden to P-behave (Diaz Legaspe 2018: 249–50; see also Ashwell 2016: 234–39). I think this is a good answer on behalf of Ashwell but I am not so sure whether it is enough to defend the counterpart assumption. Specifically, it does not support the idea that Application Thesis is entailed by Reference Thesis. Rather, as Ashwell explains, the idea of normative constraint undermines the counterpart assumption.

The second advantage is that, as Diaz Legaspe points out, Ashwell’s normative constraint does not explain derogation. In this case, my simple explanation for the entitlement to protest is beginning to look rather plausible. Every woman is entitled to protest the use of ‘slut’ because it refers to every woman. Every woman either actively behaves promiscuously or is disposed to promiscuity.

4.2 Stereotypes and the Expressive Linguistic Test

In this section, I argue that derogation stems from the stereotypes associated with the target group. Undoubtedly, there are normative and discriminatory constraints on

gender and on demographic groups but they are not the source of derogation. Diaz Legaspe agrees with Ashwell that the derogation associated with gendered slurs is related to the normative constraints on the target but she also offers a very insightful analysis of the role of stereotypes in derogation. She says, ‘More often than not discriminatory frameworks comprise stereotypical representations of the targeted class’ (2018: 251). In her view, stereotypes are always inaccurate misrepresentations because they represent the target group as a homogenous group in that the features and behavior of individuals are isolated and taken out of context and assigned to all members of the stereotyped group. ‘As a result, the public perception of members of it is reduced to a bundle of features that are thought of as natural and given’, she adds (251). Chike Jeffers makes a very similar point by saying that stereotypes take away the target’s individuality and reduce the members of the target group to caricatures. He goes on to point out that the caricatures involve false essentialist assumptions which, to me, highlights the role of essentialist generalizations in the formation of stereotypes (Glasgow et al. 2019: 69–70).

Diaz Legaspe also goes on to point out that stereotypes as inaccurate public representations are most likely a disorganized bundle of features. Hence, the grasp of stereotypes can be incomplete and might even appear idiosyncratic: ‘if a stereotype attributes features F_1 , F_2 , and F_3 to [a group], someone might grasp just F_1 and F_2 , and some others, just F_2 and F_3 .’² (Diaz Legaspe 2018: 251). With the current view, the use of slurs as a discriminatory practice is often based on stereotypes. Stereotypes are unwarranted generalizations because they reduce the target class to nothing but criminals, lazy, vulgar, and so on. The derogation associated with slurs is based on stereotypes. The meaning of slurs involve stereotypes that turn out to be extensionally trivial. The rules for ‘Boche’ assign the German people as the reference of ‘Boche’. This is because the generic operator does not have an extensional interpretation. Nevertheless, the stereotype is inferentially relevant as it allows to attribute cruelty to the German people.

At this point, an advocate of an expressive view of slurs could quite rightly ask: ‘Yes, often the use of slurs is based on a stereotypical conception but what about the instances in which the use is not based on stereotypes?’ According to expressivism, derogation is not based on any specific stereotypes. Rather, derogation is based on an expression of a negative attitude toward the target. For example, the meaning of ‘Boche’ is something like ‘Boo the Germans!’ The source of derogation is a nonpropositional attitude (see Jeshion 2013a: 307–35 and 2013b: 314–29; Potts 2007: 176–79; and McCready 2010: 5–11). This is a definite advantage of expressivism. It can explain the cases in which usage is not based on stereotypes and it can also cope with the cases in which stereotypes are involved. Expressivists can say that, in these cases, the use of a slur is still an expression of hostility toward the target and often there is a reason for the hostility. The stereotypical conception is an obvious candidate for that reason.

² Christopher Hom accurately points out that the racist social institutions ground the meanings of slurs (Hom 2008: 430–31). To me, this echoes the Putnamian division of linguistic labor. So if you want the exact stereotypical content of a given slur, consult the experts in racist language—the racist institutions.

The speaker is hostile toward the Germans *because* they are cruel or toward the French *because* they are vulgar. But an expressivist says that this is beside the point because an account of the meaning of slurs need not include the racists' personal and often complex reasons for their racism.

However, Adam Sennet and David Copp argue in a 2019 paper that, in many cases, the stereotypes are much more closely connected to the meaning of slurs. This is apparent in many slurring verbs. In the verb 'to gyp', for example, derogation very obviously derives from the stereotypical conception that associates swindling with Romani people and in the expression 'to throw a Paddy', derogation derives from the stereotypical conception that the Irish are prone to tantrums. At the same time, Sennet and Copp acknowledge that the competency in the use of 'Frog' need not include any knowledge of stereotypes. The xenophobe can simply vent his, perhaps irrational, dislike of the French and yet be a competent user of 'Frog'. Sennet and Copp conclude that the prospect for a unifying view concerning the source of derogation looks rather bleak (Sennet and Copp 2019: 1–22). I do not share their pessimism. Specifically, I do not see these as mutually exclusive alternatives. Rather, I think there is a distinction to be made between the meaning of slurs and competency in their usage. In order to be competent, the speakers need to know only that 'Frog' allows them to express hostility toward the French. However, it seems to me that the linguistic fact that 'Frog' enables to express the speaker's general hostility toward the French stems from the meaning of 'Frog', which can still include the attribution of a negative stereotype. This gives the advantage back to the stereotype view. It can explain the cases in which stereotypes are not involved and it can explain the cases in which stereotypes are obviously involved. Expressivism cannot do the latter. This leads to a situation where the derogation associated with ethnic and racial slurs stems from stereotypes but they can also be used expressively, to convey hostility toward the target. Here Reference Thesis is the only constraint on the meaning and the use of slurs. In other words, 'Boche' is a bad word for the Germans. The expressive use, then, relies on the counterpart assumption. It seems to me that this provides a test to check whether the use of a slur relies on the counterpart assumption. Specifically, this test can be used to see if the counterpart assumption applies to gendered slurs. (Ashwell also points out that expressivism has traditionally relied on the counterpart assumption. She says that according to expressivism, 'a slur derogates in virtue of the use of the slur involving expression of a problematic attitude—perhaps something like contempt—toward the neutrally definable target group' [2016: 231].)

First, let us see how this works with ethnic and racial slurs. Consider Angela Merkel announcing new restrictions to the EU budget on TV. A populist leader of another EU country then sees Merkel on TV and doesn't like the restrictions at all and exclaims, 'That Kraut!' It seems to me that this is a felicitous utterance. Yet, it has very little to do with any stereotypical conception. Rather, the slur is an expression of hostility toward Merkel on the basis that she is German. Let us then apply the same test to gendered slurs. Again, consider Angela Merkel on TV announcing that Germany welcomes refugees. A member of a domestic far-right movement then sees Merkel on TV and the racist yells 'That bitch!' Again, this

seems to be a perfectly felicitous statement. It conveys the speaker's hostile attitude toward Merkel even though the hostility is very unlikely to be based on Merkel's boisterous or self-concerned attitude. More likely, the bigot is hostile toward Merkel because she is not self-concerned enough. That is, Merkel is not taking care of her own people but inviting all sorts of people into the country. In any case, this test does support the idea that gendered slurs can be used expressively, which again supports the idea that the meaning of gendered slurs is based on the counterpart assumption. Admittedly, there are gendered slurs that are not so fitting in this scenario. For example, 'That slut!' would sound rather odd. However, even if it does, it does not make it infelicitous. Moreover, this happens with ethnic and racial slurs too. A racist might be willing to apply a slur to one member of the demographic group but might be hesitant to apply it to another member of the same group. For example, even if a racist has no problems applying the N-word to a convicted criminal, the racist might think twice before applying it to Barack Obama. Again, this can be explained with the generic component. Even though there are individuals to whom a given slur cannot be applied, slurs are still coextensional with their counterparts.

4.3 Unifying View, Contrasting View, and Dictionaries

My conclusion is that the expressive linguistic test shows that there are no significant differences between ethnic and racial slurs and gendered slurs. The meanings of ethnic and racial slurs and the meanings of gendered slurs are both based on the counterpart assumption. The result is a unifying theory of slurs just like Ashwell's view. So far, I have assumed that a unifying view is always preferable to a contrasting view like Diaz Legaspe's. She argues that the semantics of ethnic and racial slurs and gendered slurs are different. The meaning of ethnic and racial slurs is based on Reference Thesis and the meaning of gendered slurs is based on Restricted Reference. Interestingly, dictionaries support Diaz Legaspe's contrasting view. Dictionary entries equate ethnic and racial slurs with their neutral counterparts but highlight that they are derogatory.³ Gendered slurs like 'slut' and 'bitch', on the other hand, are defined in the restricted sense: 'a woman who. . .'.⁴ This is indeed interesting and problematizes unifying views like mine and Ashwell's. The problem for my view is that the dictionaries do not mention any stereotypes in their entries for ethnic and racial slurs. They merely highlight that the words are derogatory. If the derogation associated with 'Boche' is based on the

³ *Boche*: Derogatory, slang, a German, especially a German soldier (Online Collins English Dictionary); German, usually disparaging (Online Merriam-Webster). *Limey*: an insulting word for a British person (Online Cambridge Dictionary); some Americans refer to British people as limeys. Some people consider this use offensive (Online Collins Dictionary). *Frog*: an insulting word for a French person (Online Cambridge Dictionary); frog is sometimes used to refer to French people, offensive (Online Collins Dictionary).

⁴ *Slut*: a woman who has sexual relationships with a lot of men without any emotional involvement (Online Cambridge Dictionary); slut is an insulting word for a woman who has a lot of sexual partners (Online Collins Dictionary). *Bitch*: an unkind or unpleasant woman (Online Cambridge Dictionary); a malicious, spiteful, or overbearing woman; informal and often offensive (Online Merriam-Webster). *Sissy*: an effeminate man or boy; informal, disparaging (Online Merriam-Webster).

stereotypical conception that the Germans are cruel, shouldn't the dictionaries then say that?

To begin the answer, let us go back to Diaz Legaspe's point about stereotypes. She says that stereotypes are public representations that include disorganized bundles of features. Concerning ethnic and racial slurs the emphasis is on the disorganisation. From these disorganized bundles of features of F_1 , F_2 , F_3 , and so on, one speaker might grasp only F_1 and another only F_3 . However, I am convinced that if a thorough investigation were conducted, a pattern would emerge. The investigation would show that, say, F_2 is most commonly associated with ethnic or racial group G . Understandably, lexicographers have not conducted this kind of research. It would be rather laborious and, ultimately, unnecessary. The expressive linguistic test shows that the expressive aspect is an important aspect of the meaning and the use of slurs. So the emphasis on the expressive derogatory nature of slurs in dictionaries is sufficient to capture their meaning. In contrast, for some reason the stereotypes associated with gendered slurs are not disorganized in a similar fashion. Rather, the stereotypes associated with gendered slurs seems to be relatively stable as the dictionary entries show. It would be interesting to know why this is the case but I think that question is beyond the scope of this paper. The main point here is that even though lexicographers do not know the most frequent stereotype associated with ethnic or racial group G , that does not mean that there does not exist one.

If this is a sufficient explanation for the inconsistency between unifying views and dictionaries, then it turns out that a unifying view is preferable to a contrasting one. This then means that Ashwell's view and my generic view are preferable to Diaz Legaspe's view. Given the reasoning in 4.1 and 4.2, the generic view is more viable of the unifying views discussed here. The generic view yields an explanation for expressive use of gendered slurs that Ashwell's view does not. The generic view also yields a simple and an effective explanation for derogation, another area where Ashwell's view falls short.

Conclusion

The generic view is a unifying view as it defends the idea that all slurs, whether ethnic and racial or gendered, are coextensional with their counterparts. My central claim is that the evidence supporting the subset argument does not lead to the rejection of the counterpart assumption. Rather, it is equally plausible that the evidence suggests that the meaning of all slurs involves a generic element. The subset argument gets its thrust from the idea that, among the semantic facts concerning gendered slurs, there is an extra constraint—the P-behavior—that then narrows down the reference to a subset of women. According to my generic view, this extra feature is not unique to gendered slurs. Ethnic and racial slurs also add a negative stereotype to the target. This stereotype is attributed to the target in a generic way. As a result of this genericity, the application of slurs allows counterexamples even with ethnic and racial slurs. At the same time, the counterpart assumption holds.

The conclusion is that the generic view is preferable to both Ashwell's view and to Diaz Legaspe's view. Diaz Legaspe sets out to defend the counterpart assumption but

it remains a bit unclear how successful her attempt is. Strictly speaking, in her view, the meaning of gendered slurs is not based on the counterpart assumption. Moreover, according to her contrasting view, the meaning of gendered slurs is different from ethnic and racial slurs. And while Ashwell does propose a unifying view of slurs, hers is based on normative constraint and, therefore, fails to explain the expressive use of gendered slurs. At the same time, it fails to account for the full derogatory force of slurs, as Diaz Legaspe aptly points out.

PASI VALTONEN 

TAMPERE UNIVERSITY

pasi.valtonen@tuni.fi

References

- Ashwell, Lauren. (2016) 'Gendered Slurs'. *Social Theory and Practise*, 42, 228–39.
- Boghossian, Paul, and Timothy Williamson. (2003) 'Blind Reasoning'. *Proceedings of the Aristotelian Society*, 77, 225–93.
- Diaz Legaspe, Justina. (2018) 'Normalizing Slurs and Out-Group Slurs: The Case of Referential Restriction'. *Analytic Philosophy*, 59, 234–55.
- Dummett, Michael. (1973) *Frege: Philosophy of Language*. New York: Harper & Row.
- Glasgow, Joshua, Sally Haslanger, Chike Jeffers, and Quayhawn Spencer. (2019) *What Is Race?* Oxford: Oxford University Press.
- Hom, Christopher. (2008) 'The Semantics of Racial Epithets'. *Journal of Philosophy*, 105, 416–40.
- Hom, Christopher, and Robert May. (2013) 'Moral and Semantic Innocence'. *Analytic Philosophy*, 54, 293–313.
- Jeshion, Robin. (2013a) 'Expressivism and the Offensiveness of Slurs'. *Philosophical Perspectives*, 27, 307–35.
- Jeshion, Robin. (2013b) 'Slurs and Stereotypes'. *Analytic Philosophy*, 54, 314–29.
- Leslie, Sarah-Jane. (2007) 'Generics and the Structure of the Mind'. *Philosophical Perspectives*, 21, 375–403.
- Leslie, Sarah-Jane. (2008) 'Generics: Cognition and Acquisition'. *Philosophical Review*, 117, 1–47.
- Leslie, Sarah-Jane. (2017) 'The Original Sin of Cognition: Fear, Prejudice and Generalisation'. *Journal of Philosophy*, 8, 393–421.
- McCready, Elin. (2010) 'Varieties of Conventional Implicature'. *Semantics and Pragmatics*, 3, 1–57.
- Medin, Douglas, and Andrew Ortony. (1989) 'Psychological Essentialism'. In S. Vosniadou and A. Ortony (eds.), *Similarity and Analogical Reasoning* (Cambridge: Cambridge University Press), 179–95.
- Potts, Christopher. (2007) 'The Expressive Dimension'. *Theoretical Linguistics*, 33, 165–98.
- Rumfitt, Ian. (2000) "'Yes" and "No"'. *Mind*, 109, 781–824.
- Sennet, Adam, and David Copp. (2019) 'Pejorative Verbs and the Prospect for a Unifying Theory of Slurs'. *Analytic Philosophy*, 61, 130–51. <https://doi.org/10.1111/phib.12160>.
- Valtonen, Pasi. (2021) 'Generic Inferential Rules for Slurs: Dummett and Williamson on Ethnic Pejoratives'. *Synthese*, 198, 6533–51.
- Williamson, Timothy. (2009) 'Reference, Inference, and the Semantics of Pejoratives'. In J. Almog and P. Leonardi (eds.), *The Philosophy of David Kaplan* (Oxford: Oxford University Press), 137–58.