

This is an unpublished manuscript. For latest draft, please email jennifnh@usc.edu.

Busting the Ghost of Neutral Counterparts

Jennifer Foster

University of Southern California

Abstract

Philosophers have almost universally assumed that some highly general semantic relationship obtains between slurs and so-called “neutral counterpart” terms. This assumption has been fleshed out in different ways. On all extant accounts, however, it implies an unmotivated distinction between paradigmatic slur/“neutral counterpart” pairs and many pairs that theorists haven’t considered, including ‘chick flick’/‘romantic comedy’, ‘stoner’/‘cannabis user’, and ‘liberal’/‘libtard’. For pairs like these, the most intuitive theory of the target relationship involves *overlap*—both in (presumed) extension and associated stereotypes. Since (I argue) we have no good reason to distinguish pairs like ‘chick flick’/‘romantic comedy’ from paradigmatic slur/“neutral counterpart” pairs, we have good reason to accept an overlap thesis about those pairs, too. An overlap thesis can accommodate the intuitions behind more orthodox views of slurs. It also paves the way for a more sophisticated understanding of ordinary bigotry.

Keywords: slurs, neutral counterparts, overlap, derogatives, derogatory classifiers, conceptual overlap, prototype theory

§1 Introduction

The hit film *Ghostbusters* was remade in 2016, to the excitement of some and the chagrin of others. This new version shared much of its plot with the 1984 original, but with one glaring exception: its cast was *woman-led*. This fact sparked record-high levels of frustration and disappointment.¹ *Ghostbusters*, many scorned, had been made a “chick flick.”

‘Chick flick’ is a derogatory expression.² It is used to demean, or otherwise diminish the value of, the things it is applied to. Thus the husband dismisses his wife’s proposal in *Date Night*:

Date Night: A husband and wife are choosing which movie to see. The wife proposes they see “that new one with Julia Roberts.” The husband scoffs. “Why would I want to see that? *That’s a chick flick.*”

¹The trailer for the new *Ghostbusters* was “disliked” on YouTube over 600,000 times before the film was even released. It had thereby become the site’s “most disliked movie trailer ever” (Time).

²See, e.g., “Ghostbusters Director Says ‘chick flick’ Is a ‘Derogatory’ Term”.

The husband believes that films he calls ‘chick flicks’ are not worth seeing. Why? Intuitively, it is because he believes that “chick flicks” have *certain characteristic features* which he considers disvaluable—*viz.*, features which make them *the sort of thing women like*.

Probably, he feels the same way about romantic comedies. Though ‘romantic comedy’ is not derogatory in the same way ‘chick flick’ is, competent speakers know that the two expressions are related. In particular, they know that the expressions are associated with *many of the same stereotypes*, and so are applied to many of the same things.

Yet competent *users* of the expressions do not think that *all and only* chick flicks are romantic comedies.³ While many people said (1) upon the release of the new *Ghostbusters*, presumably no one said (2):

- (1) The new *Ghostbusters* is a chick flick.
- (2) The new *Ghostbusters* is a romantic comedy.

Indeed, though competent users believe that chick flicks and romantic comedies have quite a lot in common, they accept that cases like (3) and (4), at least in principle, can obtain:⁴

- (3) The new *Ghostbusters* is chick flick, but it isn’t a romantic comedy.
- (4) *Silver Linings Playbook* is a romantic comedy, but it isn’t a chick flick.

My purpose in this paper is not to theorize about the terms ‘chick flick’ and ‘romantic comedy’ My target is rather is a much broader class of expressions, including those commonly known as slurs. Slurs (such as ‘k*ke’)⁵ have been standardly assumed to bear a distinctive semantic relationship to certain *non-pejorative* expressions (such as ‘Jew’ or ‘Jewish’). It has been standardly assumed, moreover, that if a speaker recognizes this semantic relationship, and recognizes that the slur is pejorative, then she will be *competent* with the slur. By contrast, knowledge of stereotypes has standardly (and especially since ?) been considered unnecessary for basic competence.

In contrast, I will argue that just as competence with ‘chick flick’ requires familiarity with stereotypes, so too does competence with expressions like ‘k*ke’. Moreover, I will argue, we have no good reason to assume that sentences like (5) and (6) are, *semantically*, any different from (3) and (4):

- (5) He’s a Jew, but he’s not a kike.⁶
- (6) He’s a kike, but he isn’t a Jew.⁷

³Throughout the paper I will distinguish between “competent speakers” and “competent users” in order to allow for speakers who are competent with how an expression is used, but do not themselves use it, or use it anymore.

⁴There is, for example, an entire Reddit discussion dedicated to the difference between ‘chick flick’ and ‘romantic comedy’.

⁵Throughout this paper, I have taken several steps in order to minimize how many times such expressions are mentioned. The worst ones are transcribed with an asterisk (*), *except* those that appear in in-set example sentences. Most of these example sentences come verbatim from the cited online sources. Some of them, however, have been edited for length or clarity to minimize the reproduction of slurs.

⁶Twitter, 22 Mar 2018.

⁷Twitter, 13 Sept 2013.

Expressions like ‘k*ke’ and ‘Jew’, we should think, bear the *same basic relationship* as ‘chick flick’ and ‘romantic comedy’. But if this is right, then we should think that nearly all extant theories of slurs are mistaken.

§2 Setup

There are certain expressions that everyone agrees are slurs. Beyond such paradigmatic examples, however, it is controversial how broadly or narrowly the term ‘slur’ should be defined. Where should we draw the line, if anywhere, between slurs and other pejorative expressions?

Traditionally, theorists have assumed that the answer to this question is at least partially semantic. This, ultimately, is an assumption that I want to challenge. But some philosophers are inclined to restrict the technical meaning of ‘slur’ for reasons that have nothing to do with semantics. Geoff ?, likewise skeptical of a sharp semantic definition, highlights a few. For example, we might think that ‘slur’ should capture only those expressions used to derogate members of groups unwarrantingly, as on the basis of traits beyond their control.⁸ Alternatively, we might want to reserve ‘slur’ for only those expressions used to derogate members of protected classes, or groups which are systemically oppressed.⁹ There is no settled answer to these questions, and, for this paper at least, I want to remain maximally neutral about them.¹⁰ My interest, rather, is in one of the most “settled” claims about slurs—viz., that expressions like ‘k*ke’ and ‘ch*nk’ bear a special semantic relationship to non-slur expressions like ‘Jew’ and ‘Chinese’. This claim is so standardly assumed, in fact, that bringing it into view will require some setup—and some nonstandard terminology.

This means I will be setting aside the word ‘slur’ for the majority of what follows. I will instead talk of what I’ll call *derogatory classifiers*, or “DCs.” DCs include all of the expressions that theorists typically call slurs, such as ‘k*ke’, ‘ch*nk’, ‘n*gger’, and ‘f*ggot’. But they also include many expressions that theorists generally do not analyze the same way, such as ‘d*ke’, ‘c*nt’, ‘slut’, ‘towelhead’, ‘commie’, ‘libtard’, ‘nazi’, ‘boomer’, ‘Bernie bro’, ‘Trumper’, ‘redneck’, ‘hillbilly’, ‘chav’,¹¹ ‘treehugger’, ‘cheesehead’, ‘townie’, ‘pillhead’, ‘junkie’, ‘alchie’, ‘stoner’, ‘fatso’, ‘ginger’, ‘nerd’, ‘geek’, ‘gamer’, ‘frat bro’, ‘trailer trash’,

⁸?, who observes: “We might speak of a word for the members of a group as derogative even if we personally think they merit derogation. . . . But most of us would demur from calling either word a slur, since we feel the groups have it coming” (239). This view is explicitly suggested by ?.

⁹In this vein ? writes, “a derogative word qualifies as a slur only when it disparages people on the basis of properties such as race, religion, ethnic or geographical origin, gender, sexual orientation, or sometimes political ideology—the deep fatalities that have historically been the focus of discrimination or social antagonisms that we see as rents in the fabric of civil society” (239). It is in this sense, I take it, that many people deny that terms like ‘nazi’ and ‘cracker’ are slurs. Similar questions have arisen in the recent ‘Is ‘TERF’ a slur’ debate ??.

¹⁰For a nice discussion, see ?.

¹¹The derogatory expression ‘chav’ is commonly used in the UK to refer to “a young person of a type characterized by brash and loutish behaviour and the wearing of designer-style clothes (esp. sportswear); usually with connotations of a low social status” (OED). Thanks to Zoe Johnson King for introducing me to this expression.

‘traphouse’, ‘McMansion’, ‘dive bar’, ‘fleabag motel’, ‘man cave’, ‘mom jeans’, ‘dad joke’, and ‘chick flick’.¹²

When presented with such examples, competent speakers find it easy to generate more. And intuitively, DCs form a broad linguistic class. I wish not to beg questions, however, about what makes this class *unified*. In the first place, I wish not to assume anything general about DCs’ status as “slurs”. Obviously, expressions like ‘frat bro’, ‘boomer’, ‘McMansion’, and ‘dad joke’ are not systematically oppressive; and however offensive they may be to certain communities of speakers, they carry nothing of the hideous force of the n-word. Whether and when an expression should be called a slur are considerations as heavy as they are fraught, and it is an advantage of DC-talk that it does not (or at least need not) carry the same weight.

More importantly, though, I wish not to beg questions about DCs’ semantic status. Broadly, DCs are nominalized, variably pejorative¹³ (American) English expressions for categories of persons or things. I assume that they include all (but not only) expressions typically taken to derogate “on the basis of such things as race, ethnicity, nationality, class, religion, ideology, gender, and sexual orientation” (? , p.60). But in saying that they are “pejorative” or “derogative”, I mean only that DCs have at least some pejorative or derogative *uses* in certain linguistic communities.¹⁴ I leave it open whether those uses are explained by something general about DCs’ meaning, or by something general about their socio-political contexts of utterance. Indeed, it is the principal advantage of DC-talk that it permits a theory-neutral way of asking the central semantic question at issue in this paper. DC-talk is neutral, in the first place, about whether the derogatory force of prototypical slurs (like

¹²? similarly claims that paradigmatic slurs belong to a broader class of expressions, which he calls “derogatives”, and are distinguished pragmatically rather than semantically. But while I agree with Nunberg that (paradigmatic) slurs have much more in common with, e.g., expressions like ‘flack’, ‘techie’ and ‘treehugger’ than most theorists recognize (pp. 239-241), I disagree with him about their semantics. And while I, too, want to cast a much wider net, linguistically, than is usually cast in discussions of slurs, I am unsure that Nunberg and I want to cast the *same* wide net. For these reasons, and at the risk of multiplying theoretical terms beyond necessity, I have introduced my target expressions by a different name.

¹³Derogatory variation—i.e., that pejorative expressions can vary in their derogatory force or intensity—has been long-observed in the philosophical literature on slurs (? , p.426; ? , p.233; ? , p. 241). Little attention has been paid, however, to expressions on the lower, or “milder,” end of the spectrum. Many of the expressions I cite as DCs are near (or nearer) that end. These include expressions, like ‘nerd’ and ‘townie’, which are generally acknowledged to be pejorative, but only mildly so; expressions like ‘c*nt’ and ‘chav’, which are generally acknowledged to be pejorative in some dialects of English, but not others; and expressions like ‘stoner’, ‘gamer’, and ‘man cave’, which are not universally acknowledged to be pejorative, but are clearly used and interpreted pejoratively by certain subcommunities of speakers.

¹⁴Theorists often distinguish between “group slurs,” like ‘k*ke’ and ‘commie’, and “personal slurring terms,” like ‘fatso’, and ‘jerk’, and ‘asshole’ ????. I leave it open whether DCs include pejoratives like these, which are not (obviously) restricted to particular groups or categories. I am not convinced myself, however, that there is a clear-cut distinction here. Pejoratives like ‘redneck’, ‘nerd’, and ‘Bernie bro’ are used to classify relatively narrowly circumscribed groups, even if those groups are characterized by individual choices or actions. The same seems true of the expressions ‘dumbass’, ‘idiot’, ‘moron’, and ‘r*tard’, where these expressions are clearly related but admit of significant derogatory variation. I’m inclined to think that expressions like ‘asshole’ and ‘jerk’, which ? calls “all purpose putdowns”, are part of another kind of “continuum,” with expressions like ‘douchebag’, ‘dirtbag’, and ‘skeezeball’ falling somewhere between them and more narrowly-defined DCs. Notably, Jeshion herself expresses doubt that there is “any real *semantic* difference” between expressions like ‘f*ggot’, on the one hand, and ‘fatso’ on the other ?, p. 236.

‘k*ke’) is semantically explained; and among views on which it is semantically explained, DC-talk is neutral about what explains it.

As I’ve said, ordinary English speakers have a sense that DCs “go together”. They may disagree about whether each particular expression on my list, above, is actually a derogatory classifier term; but if asked to generate their own list of DCs, they easily could. Such speakers, moreover, are liable to notice a trend among many of the DCs on their list. In particular, they are liable to notice that many DCs may be intuitively “paired up” with other category expressions that are not (obviously) pejorative in their typical uses. I’ll call these more “neutral” group or category expressions *non-pejorative associates*, or “NPAs”:

Derogatory Classifier (DC)	Non-Pejorative Associate (NPA)	Derogatory Classifier (DC)	Non-Pejorative Associate (NPA)
‘k*ke’	‘Jew’	‘cheesehead’	‘person from Wisconsin’
‘n*gger’	‘black’	‘townie’	‘[town] native’
‘ch*nk’	‘Chinese’	‘junkie’	‘(heroin) drug addict’
‘f*ggot’	‘gay’	‘alchie’	‘alcoholic’
‘d*ke’	‘lesbian’	‘stoner’	‘weed smoker’/‘cannabis user’
‘c*nt’	‘woman’	‘pillhead’	‘prescription drug addict’
‘towelhead’	‘arab’	‘fatso’	‘fat person’
‘wetback’	‘(Mexican) immigrant’	‘ginger’	‘redhead’
‘commie’	‘communist’	‘gamer’	‘person who plays video games’
‘libtard’	‘liberal’	‘frat bro’	‘fraternity brother’
‘Bernie bro’	‘Bernie Sanders supporter’	‘mom jeans’	‘jeans worn by moms’
‘trailer trash’	‘poor person’	‘dad joke’	‘joke told by dads’
‘boomer’	‘(American) person born 1946-1964’	‘man cave’	‘den’/‘garage’/‘basement’
‘treehugger’	‘environmentalist’	‘chick flick’	‘romantic comedy’

Philosophers have long observed that paradigmatic slurs like ‘k*ke’ and ‘f*ggot’ seem to come in pairs with expressions like ‘Jew’ and ‘gay’.¹⁵ Standardly, these expressions are called “neutral counterparts.”¹⁶ Indeed, slur/neutral counterpart pairs like ‘k*ke’/‘Jew’ have been central to most theorizing about slurs.

As we will see, most theorists would deny that all of the DC/NPA pairs in the table above are genuine “slur/neutral counterpart” pairs.¹⁷ Still, I assume that NPAs include all of the expressions theorists typically called “neutral counterpart” terms. Additionally, I assume that only some DCs have NPAs: some, like ‘chav’, ‘nerd’, and ‘redneck’, intuitively do not.

The list of DC/NPA pairs above is easy to generate. Indeed, I suspect that the reader, too, having detected a pattern, will find it easy to keep going. My default hypothesis, then, is that DC/NPA pairs have something basic in common. In particular, my default hypothesis is that DC/NPA pairs exhibit the *same basic relationship*. If this hypothesis is true, then we should expect to find systematic behavior across DC/NPA pairs. That this is precisely what we do find is my central contention.

Whatever the relationship is between DCs and NPAs, it is partially contrastive. This is to some extent true by definition, insofar as DCs are typically pejorative expressions,

¹⁵Many philosophers have offered ‘homosexual’ as the “neutral counterpart” of ‘f*ggot’. But in many linguistic communities (including my own) ‘homosexual’ has become increasingly socially marked.

¹⁶? and ? call them “nonpejorative correlates.”

¹⁷I discuss more orthodox views of “neutral counterparts” in §§6-7.

and NPAs typically are not. But beyond this difference in (typical) pejorative force, DCs and NPAs are also used in *descriptively* contrastive ways. The following kinds of normative judgments and commands, for example, are made by competent users all the time:

- (7) I have nothing against Jews. I do not like kikes.¹⁸
- (8) It's fine being black but don't be a nigger everywhere you go.¹⁹
- (9) I'll support the shit out of women all day long. We're fabulous. However..it's not OK to be a cunt just because you have one.²⁰
- (10) Be a lesbian & not a dyke, theres a huge difference...²¹
- (11) It's okay to be gay, it's not okay to be a faggot tho.²²
- (12) It's okay to be liberal, but it's not okay to be a libtard.²³
- (13) I'm OK with people who support Bernie. I do not like Bernie Bros.²⁴
- (14) Look, I'm all for the environment, but tree huggers get on my everlasting nerve. Shut UP.²⁵
- (15) I'm not a 'hippie', a 'tree-hugger' or any other stereotypical name you put to it - but if you're not an environmentalist, you're fucked up.²⁶
- (16) Hey Wisconsinians don't be a cheesehead.²⁷
- (17) I like redheads, but not gingers.²⁸
- (18) I have nothing against cannabis users, but I hate stoners.²⁹
- (19) It's okay to play video games. It is not okay to be gamer.³⁰
- (20) Even if you are a mom, it's not okay to wear mom jeans.³¹
- (21) Act like a fraternity brother, and not a frat bro.³²
- (22) Craft rooms, hobby rooms, office rooms, and dens are fine. But I draw the line at man caves.³³
- (23) I like romantic comedies. Not chick flicks. There's a huge difference.³⁴

An adequate theory of the DC/NPA relationship, then, should explain not only why competent users intuitively *associate* particular DCs and NPAs, but also why they frequently *differentially evaluate them*.

¹⁸Nadiya Savchenko, Ukrainian Parliament Member, 29 Mar 2017.

¹⁹Twitter, 23 May 2019.

²⁰Twitter, 9 Mar 2018.

²¹Twitter, 28 Mar 2011.

²²Twitter, 13 Nov 2016.

²³Twitter, 2 Sep 2018.

²⁴Twitter, 20 May 2016.

²⁵Twitter, 24 Dec 2011.

²⁶Twitter, 11 May 2015.

²⁷Twitter, 30 Mar 2016.

²⁸Twitter, 29 Aug 2012.

²⁹Twitter, 26 Mar 2019.

³⁰Twitter, 25 Nov 2018.


³¹Twitter, 27 Dec 2012.

³²Twitter, 23 Mar 2013.


³³Twitter, 2 Sep 2010.

³⁴Twitter, 27 Jan 2012.


A key observation of this paper is that, for ‘chick flick’/‘romantic comedy’, the most plausible theory of this relationship intuitively involves *overlap*. Competent users of these expressions do not believe that all or merely some romantic comedies are chick flicks, but that *most* romantic comedies are chick flicks—and vice versa. In this sense, I will say that the target relationship involves *presumed extensional overlap* (*pE*-overlap). Competent users believe (perhaps mistakenly) that both ‘chick flick’ and ‘romantic comedy’ have non-empty extensions; and that those extensions contain most (if not all) of exactly the same things.³⁵



Restricted Extension. All chick flicks are romantic comedies, but not all romantic comedies are chick flicks.



Overlap. Most but not all chick flicks are romantic comedies, and most but not all romantic comedies are chick flicks.



Coextension. All and only chick flicks are romantic comedies, and all and only romantic comedies are chick flicks.

But competent users also associate ‘romantic comedy’ with most (if not all) of the stereotypes, beliefs, and attitudes they associate with ‘chick flick’. In this sense, the target relationship involves a substantial degree of what I’ll call *conceptual overlap* (*C*-overlap).³⁶ Broadly, then, the relationship between ‘chick flick’ and ‘romantic comedy’ involves *pE*- and *C*-overlap. Call this an *overlap thesis* about the DC/NPA pair.

My core argument, then, is simple. Since **(P1)** we have good reason to accept an overlap thesis about ‘chick flick’/‘romantic comedy’, and since **(P2)** we have no good reason to distinguish ‘chick flick’/‘romantic comedy’, semantically, from other DC/NPA pairs, then **(C)** we have good reason to accept an overlap thesis about DC/NPA pairs in general.

I provide some additional support for (P1) in the next section (§3), though will I take its plausibility largely for granted. Most of the paper is an argument for (P2). Importantly, (P2) is a claim about which general semantic distinctions we have reason (or not) to draw between ‘chick flick’/‘romantic comedy’ and other DC/NPA pairs. It is emphatically *not* the claim that there are, in fact, *no* deep semantic differences to capture. My goal is not to show, definitively, that the kind of overlap relationship that intuitively obtains between

³⁵This is consistent with some DCs having null extension, as has been suggested for slurs by e.g., ??, and ?

³⁶Does that mean that stereotypes are part of the *semantics* of DCs? Not necessarily. As I have described *pE*- and *C*-overlap, an “overlap view” of DCs is consistent with a range of linguistic frameworks, which may account for the role *C*-overlap differently. *Inferential pragmatic frameworks*, for example, may account for *C*-overlap wholly pragmatically. On these frameworks, lexical entries are atomic, but activate “mental files” containing associated beliefs and representations. ????. By contrast, *dual-content semantic frameworks* and *prototype semantic frameworks* may account for *C*-overlap wholly semantically. On dual-content semantic frameworks, lexical entries have an extensional structure (E-structure) that atomically determines extension and prototype or conceptual structure (C-structure) that consists of associated beliefs or representations ????. Prototype semantic frameworks do not distinguish between E- and C-structures, holding instead that lexical entries encode highly structured prototype concepts governed by similarity relations ?????. A broadly inferentialist view of slurs has been given by ?, and prototype theories of slurs have been given by ?? and ?.

‘chick flick’ and ‘romantic comedy’ is the *fundamental* relationship for all DC/NPA pairs. My goal, rather, is to shift the burden to those who have assumed it is not.

Nearly all philosophers theorizing about slurs have arrived at the same (in my view, faulty) generalization—*viz.*, that DC/NPA pairs like ‘k*ke’/‘Jew’, ‘ch*nk’/‘Chinese’, and ‘f*ggot’/‘gay’ exhibit a *special* semantic relationship that distinguishes them from other DC/NPA pairs like ‘chick flick’/‘romantic comedy’, ‘stoner’/‘cannabis user’, ‘lib-tard’/‘liberal’, and even ‘c*nt’/‘woman’. Philosophers have arrived at this generalization, I believe, by focusing on far too narrow an evidence set. Accordingly, I begin my argument for (P2) by taking stock of more of the data. As I highlight in §4 and §5, there is striking *systematicity* across the DC/NPA pairs just mentioned—as well as the pairs ‘mom jeans’/‘jeans worn by moms’ and ‘dad joke’/‘joke told by dads’—in *precisely* the sorts of linguistic contexts that have motivated more orthodox views of pairs like ‘k*ke’/‘Jew’. This systematicity, *prima facie* at least, seems to belie any *special* relationship between paradigmatic slurs and their so-called “neutral counterparts.”

As I have suggested, this is a heterodox conclusion. In §6 and §7, I explain why we should think it is true. First, I show that my default hypothesis—*viz.*, that DC/NPA pairs exhibit the same basic relationship—has clearly false implications for most orthodox theories. Proponents of these theories, then, have two options: they must either (a) explain the parity between pairs like ‘k*ke’/‘Jew’ and ‘chick flick’/‘romantic comedy’ by means of some other general linguistic mechanism, or (b) postulate different semantic “kinds” of DC/NPA pairs. As I show in §7, however, both of these options are empirically unmotivated. They are also unnecessary: a basic overlap thesis for DC/NPA pairs, like the one I sketch in §3 for ‘chick flick’/‘romantic comedy’, can get us the same benefits with none of the costs.

Why, then, have so many philosophers been attracted to more orthodox theories? In Section §8, I offer a diagnosis.

§3 Overlap

‘Chick flick’ and ‘romantic comedy’, I have claimed, are related by two kinds of *overlap*. Extensionally, they are presumed by competent users to pick out many (if not all) of the same films (what I have called *pE*-overlap). And conceptually, they are associated by competent users with many of the same stereotypes, beliefs, and attitudes (what I have called *C*-overlap). This, I take it, is an intuitive thesis about ‘chick flick’/‘romantic comedy’. In this section, however, I’ll say a little more to motivate it. In particular, I’ll argue that an overlap thesis about ‘chick flick’/‘romantic comedy’ can explain exchanges like (23):

- (23) A: The new *Ghostbusters* is a chick flick.
 B: But it’s an action film, not a romantic comedy!
 A: Whatever, it’s still a chick flick.

When competent users of ‘chick flick’ and ‘romantic comedy’ witness exchanges like (23), they effortlessly understand them. But what explains this, exactly? Why should B’s utterance be a *relevant*—let alone *natural*—reply to A’s? What do romantic comedies, which A never mentioned, have to do with anything?

Intuitively, the answer is that B has *taken something for granted* about the relationship between chick flicks and romantic comedies. In particular, B seems to have presupposed


that a film is a chick flick if (and only if) it is also a romantic comedy. Thus, as he believes that *Ghostbusters* isn't a romantic comedy, he tries to correct A's application of 'chick flick' to the film. Call exchanges like (23) *DC-corrections*.

On reflection, B would likely *reject* this presupposition. If asked to provide examples like (3) and (4), he probably could:

(3) The new *Ghostbusters* is chick flick, but it isn't a romantic comedy.

(4) *Silver Linings Playbook* is a romantic comedy, but it isn't a chick flick.

Still, he would probably have to *think* about it. For when B, like most competent users, ordinarily hears³⁷ the words 'chick flick' and 'romantic comedy', cases like (3) and (4) are probably not what comes to mind. What likely comes to mind, ordinarily, are films like *Pretty Woman*, *Love Actually*, *Bridesmaids*, and *He's Just Not That Into You*—*viz.*, films that are *both* chick flicks and romantic comedies, as far as B is concerned. For most competent users, prototypical examples of chick flicks are *also* prototypical examples of romantic comedies. So while cases like (3) and (4) are *possible*, they are far from typical. They obtain, in my Venn diagram metaphor, only at the *margins*.



Overlap. Most but not all chick flicks are romantic comedies, and most but not all romantic comedies are chick flicks.

Such beliefs about prototypical examples (*C*-overlap) inform competent users' *default expectations* about exchanges like (23). In particular, they give rise to a *practical presumption against marginal cases*.

As competent users, both A and B believe that sentences like (3) and (4) can be true in principle. They believe that, extensionally, 'chick flick' and 'romantic comedy' can at least conceivably come apart. But A and B also believe that this possibility *rarely matters in practice*. They know that, as far as they and most other competent users are concerned, exceptions to the rule "all and only chick flicks are romantic comedies" typically do not obtain. Thus *C*-overlap sets their default expectations about not only which movies will be "chick flicks" and "romantic comedies," but which *conversational moves* they will be able to make with each other. Their shared presumption against marginal cases is part of the common ground.

This explains why B replies the way he does in (23). B *takes it for granted* that when someone uses the word 'chick flick', they also mean 'romantic comedy' (and vice versa). If

³⁷For the sake of succinctness, I will describe conversational participants as "speakers" and "hearers" throughout. I recognize, however, that not all speech is hearable—and not everyone who communicates hears. My intention, then, is that all instances of "speakers" and "hearers" be read as shorthand for "communicators who speak, write, type, and sign" and "communicators who hear, read, feel, and see," respectively. I regret not finding a less ableist convention.

A had wanted to *cancel* this presumption against marginal cases, he could have asserted (3):

(3) The new *Ghostbusters* is a chick flick, but it isn't a romantic comedy.

But A has not asserted (3); he has asserted (24):

(24) The new *Ghostbusters* is a chick flick.

And this, naturally, raises the possibility of (25) to salience—which B rejects:

(25) The new *Ghostbusters* is a romantic comedy.

Hence his attempted correction: “But it’s an action film, not a romantic comedy!”

Intuitively, B takes A to have made some kind of *mistake*. Perhaps he believes that A has made a *semantic* mistake—*viz.*, that A has applied ‘chick flick’ to a film that (by B’s own lights) is not in the term’s extension. But we need not assume this. B’s complaint against A, *qua* fellow competent user, could just as easily be *pragmatic*. By asserting (24) when he could have asserted (3), A has omitted to cancel the otherwise operative presumption that all and only chick flicks are romantic comedies—and has thus predictably and avoidably introduced confusion into their exchange. Indeed, this explains a notable feature of DC-corrections like (23): though A intuitively *concedes* B’s rejection of the NPA —‘Whatever’— he *feliculously doubles down* on his application of the DC. B’s attempted correction, whatever mistake he takes A to have made, is in this crucial sense not conversation-stopping.

An adequate theory of relationship between ‘chick flick’ and ‘romantic comedy’ must be able to explain this datum. In particular, it must be able to explain not only why competent users take the expressions to be intuitively similar (extensionally, conceptually), but also why they sometimes apply them to different things. An overlap thesis about ‘chick flick’/‘romantic comedy’ does this. It captures competent users’ intuitions about what the two expressions mean; it explains, with some plausible assumptions about common ground, why such users would share a default presumption against marginal cases; and it predicts the right linguistic judgments about (23)—*viz.*, that while B’s reply is natural, it is not conversation-stopping.

As I will show in the next section, however, the same considerations appear to be systematically paralleled by other DC/NPA pairs. And if this is right, then we should think that an overlap thesis is also adequate for the DC/NPA relationship more generally.

§4 Data

In the previous section, I emphasized how competent users of ‘chick flick’ and ‘romantic comedy’ believe that some sentences like (3) and (4) are, or in principle could be, true:

(3) The new *Ghostbusters* is chick flick, but it isn't a romantic comedy.

(4) *Silver Linings Playbook* is a romantic comedy, but it isn't a chick flick.

In these cases, they take the DC and the NPA to come apart in *both directions*.

But *prima facie*, at least, competent users seem to use other DCs and NPAs in the same, bidirectionally divergent way. Indeed, prior to questions of semantics, there seems to be

substantial *syntactic* parity between ‘chick flick’/‘romantic comedy’ and other DC/NPA pairs in sentences like (3) and (4). Examples of such isomorphism are plentiful on Twitter—and my goal in this section is merely to showcase their range:

‘boomer’/‘(American) person born 1946-1964’

(26) Grandma is not a boomer, she’s cool.³⁸

(27) Mayo Pete is a boomer at 37. Does this help everyone understand that boomer isn’t just an age thing?³⁹

‘Bernie bro’/‘Bernie Sanders supporter’

(28) I’m still a Bernie fan. But I’m not a Bernie Bro. Do you know what I mean? I enjoy the guy. I’d like him to be President. But I’m not gonna be mad that you enjoy another candidate. And I expect the same respect in return. This is called Democracy - or so I’m told.⁴⁰

(29) Next to Misogyny in Websters is a photo of this hateful horrible Republican Bernie bro... #HatesWomen⁴¹

‘cheesehead’/‘person from Wisconsin’

(30) I’m not from Wisconsin, but I’m still a cheesehead.⁴²

(31) I’m from Wisconsin but not by any means am I cheesehead.⁴³

‘dad joke’/‘jokes told by dads’

(32) Being a dad with “actually funny jokes and not just dad jokes” might be the best feedback you can get from your kid.⁴⁴

(33) I don’t understand why Cory Booker tells so many Dad jokes when he’s not actually a Dad.⁴⁵

‘fatso’/‘fat person’

(34) There’s a difference between being fat cause you love eating but you still look hott and BEING FAT cause your just a fatso.⁴⁶

(35) im skinny but i eat alot still so im a fatso.⁴⁷

‘gamer’/‘person who plays video games’

³⁸Twitter, 14 Nov 2019.

³⁹Twitter, 19 Nov 2019.

⁴⁰Twitter, 27 Sep 2019.

⁴¹Twitter, 16 Nov 2019.

⁴²Twitter, 27 Apr 2012.

⁴³Twitter, 15 Sep 2013.

⁴⁴Twitter, 2 Apr 2019.

⁴⁵Twitter, 15 Oct 2019.

⁴⁶Twitter, 15 May 2013.

⁴⁷Twitter, 26 Jun 2012.

(36) I am a game player, for sure, but not a gamer.⁴⁸

(37) Typical gamer. Doesn't even play the game, just wanks off at the pretty girls.⁴⁹

'ginger'/'redhead'

(38) She's a redhead but not a ginger.⁵⁰

(39) She's a ginger, not the redhead ginger though.⁵¹

'junkie'/'(heroin) drug addict'

(40) Getting addicted does in fact, make you an addict! But it does not make you a junkie.⁵²

(41) If you snort cocaine from time to time your still a junkie i don't care, i don't care.⁵³

'libtard'/'liberal'

(42) I wasn't saying all liberals are libtards. They are not. But libtards do exist and they need to be called out on their bullshit.⁵⁴

(43) The irony was amaze balls and I now get it. Not all libertarians are libtards. But the ones who think their way is the only way most definitely are.⁵⁵

'mom jeans'/'jeans worn by moms'

(44) My mom doesn't wear mom jeans. She wears cool jeans.⁵⁶

(45) Obama wears mom jeans, but he isn't not a mom.⁵⁷

'stoner'/'cannabis user'

(46) Just because someone is a cannabis user doesn't mean they're a stoner.⁵⁸

(47) Just because someone isn't a cannabis user doesn't mean they're not a stoner.⁵⁹

'trailer trash'/'poor person'

(48) Trump family: Proving you don't have to be poor to be trailer trash.⁶⁰

(49) Trailer trash, anyone? And let's not forget everyone else who lives in poverty.⁶¹

⁴⁸Twitter, 23 Jun 2019.

⁴⁹Twitter, 30 May 2014.

⁵⁰Twitter, 3 Oct 2017.

⁵¹Twitter, 13 Sep 2012.

⁵²Twitter, 19 Mar 2019.

⁵³Twitter, 12 Oct 2018.

⁵⁴Twitter, 20 Aug 2018.

⁵⁵Twitter, 16 Mar 2018.

⁵⁶Twitter, 31 May 2012.

⁵⁷Obama on 'mom' jeans: 'I've been unfairly maligned'.

⁵⁸Twitter, 1 Oct 2015.

⁵⁹Twitter, 23 May 2014.

⁶⁰Twitter, 22 Nov 2017.

⁶¹Twitter, 3 Mar 2013.

This apparent pattern extends also to DC/NPA pairs more commonly discussed in the slurs literature:⁶²

- (50) Not all chinese people are “chinks”. And not all black people are “niggers”. Not all italians are “guineas”. Not all white people are “white trash”. There is a difference.⁶³
- (51) I hate those stupid Vietnamese people. Stupid chinks always torturing animals.⁶⁴
- (52) He isn’t a kike like 99% of the Jews.⁶⁵
- (53) *said about investment bankers:*
To fix the economy you have to tax the kikes on Wall Street.⁶⁶
- (54) There are gays, and then there are faggots.⁶⁷
- (55) *said about an executive ban on transgender military personnel:*
Trump won’t allow faggots in the military. [But] I say if they can do the job of a soldier, let ‘em serve.⁶⁸
- (56) I always wanted a lesbian friend, lesbian not a dyke.⁶⁹
- (57) Not all dykes are lesbians. I got a cousin who’s a dyke but she has a husband.⁷⁰
- (58) Not all women are cunts. Term is reserved for only the deserving. Men can be cunts too. Just as they can be pussies.⁷¹

The structural pattern, here, is striking—as far as it goes. But while such systematicity is *consistent* with my default hypothesis—*viz.*, that DC/NPA pairs exhibit the same basic relationship—it certainly does not entail it. Much more needs to be shown to motivate an overlap thesis (of the kind sketched in §3 for ‘chick flick’/‘romantic comedy’) for DC/NPA pairs more generally. This is my goal in the next section.

§5 Systematicity

An overlap thesis, recall, says of a given DC/NPA pair that the relevant relationship involves substantial presumed extensional overlap (*pE*-overlap) grounded in substantial conceptual overlap (*C*-overlap). Such a thesis, I have claimed, intuitively captures competent users’ intuitions about ‘chick flick’/‘romantic comedy’, and nicely explains important features of (23):

- (23) A: The new *Ghostbusters* is a chick flick.
B: But it’s an action film, not a romantic comedy!
A: Whatever, it’s still a chick flick.

⁶²Not incidentally, such uses of slurs have been largely ignored in the philosophical literature. The most notable exception is Adam ??, though ?? and ? have (to very different ends) also emphasized such uses.

⁶³Online comment, 4 July 2013.

⁶⁴Twitter, 5 June 2015.

⁶⁵Twitter, 20 November 2016

⁶⁶Twitter, 25 June 2013.

⁶⁷Twitter, 17 Nov 2017.

⁶⁸Twitter, 27 July 2017.

⁶⁹Twitter, 28 June 2013.

⁷⁰Twitter, 9 Nov 2019.

⁷¹Twitter, 11 Dec 2016.

In this section, I will show how the same considerations intuitively extend to other DC/NPA pairs. I will focus on the following DC/NPA pairs in particular:

‘mom jeans’/‘jeans worn by moms’	‘stoner’/‘cannabis user’	‘c*nt’/‘woman’
‘dad joke’/‘joke told by dads’	‘libtard’/‘liberal’	‘d*ke’/‘lesbian’

These pairs pattern with ‘chick flick’/‘romantic comedy’ in DC-corrections like (23). Moreover, this behavior supports the hypothesis that competent use of the various expressions requires familiarity with relevant associated stereotypes.

‘mom jeans’/‘jeans worn by moms’ and ‘dad joke’/‘joke told by dads’

Like ‘chick flick’, ‘mom jeans’ and ‘dad joke’ are DCs: they are nominalized group or category expressions with pejorative uses in certain linguistic communities. They also have intuitive NPAs—‘jeans worn by moms’ and ‘joke told by dads’, respectively—where the former picks out all and only the jeans worn by moms, and the latter picks out all and only the jokes told by dads. And as with ‘chick flick’/‘romantic comedy’, the relationship between these DC/NPA pairs intuitively involves *pE*- as well as *C*-overlap.

Competent users of ‘mom jeans’ believe that, of all the jeans worn by moms, many are mom jeans. Likewise, competent users of ‘dad joke’ believe that many of the jokes told by dads are dad jokes. Still, they believe that people other than moms can wear moms jeans, and that people other than dads can tell dad jokes:

- (32) Being a dad with “actually funny jokes and not just dad jokes” might be the best feedback you can get from your kid.
- (33) I don’t understand why Cory Booker tells so many dad jokes when he’s not actually a dad.
- (44) My mom doesn’t wear mom jeans. She wears cool jeans.
- (45) Obama wears mom jeans, but he’s not a mom.

Essential for understanding these sentences, moreover, is familiarity with the relevant stereotypes. For ‘mom jeans’/‘jeans worn by moms’, these include features like *loose*, *high-waisted*, and *unflattering*; and for ‘dad joke’/‘joke told by dads’ they include features like *trite*, *punny*, and *groan-inducing*. Even more basically, competent speakers must be familiar with stereotypical features associated with ‘mom’, ‘jeans’, ‘dad’, and ‘joke’—such as *being uncool/embarrassing*, *being made of denim*, *being silly with kids*, and *being intended to be funny*.

In the case of both pairs, the DCs bear an explicit *etymological* connection to their NPAs. And given this explicit connection, it is even less surprising than it was in the case of ‘chick flick’/‘romantic comedy’ that competent users might have exchanges like (59) and (60):

- (59) A: Obama wore mom jeans to the baseball game.
 B: But Obama’s not a mom!
 A: Whatever, he still wore mom jeans.
- (60) A: Ellen told dad jokes at the Oscars.
 B: But Ellen’s not a dad!
 A: Whatever, she still told dad jokes.

In both DC-corrections, B's replies are relevant, natural, and intuitively conceded by A. But just as in (23), this does not keep A from felicitously doubling down in both cases on his original claim. While B's attempted corrections pressure A to reject the NPAs, they do not pressure A to reject the DCs.

'stoner'/'cannabis user'

Though 'stoner' is less derogatory than many DCs, it has unquestionably pejorative uses.⁷² In particular, competent users associate the expression with negative stereotypes like *frequently using mind-altering drugs, sleeping too much, being lazy, being unemployed, being dirty, eating junk food, and being unintelligent*. Not incidentally, these are also stereotypes that competent users associate with the more neutral 'cannabis user', along with variant expressions like 'marijuana user' and 'weed smoker'. Competent users of 'stoner' and 'cannabis user', then, exhibit *C*-overlap.

They also exhibit *pE*-overlap: they believe that many (if not most) cannabis users are "stoners."⁷³ Still, as (46) and (47) suggest, most of these speakers are prepared to make exceptions. Much as certain romantic comedies (like *Silver Linings Playbook*) are judged "too serious," "too important," or "too complex" to be "chick flicks,"⁷⁴ certain cannabis users are judged "too productive", "too smart", or "too clean-cut" to be "stoners." Likewise, much as the 2016 *Ghostbusters*, while judged not to be a romantic comedy, was judged to be a "chick flick," many individuals judged not to be cannabis users are judged to be "stoners." Speakers who make such judgments believe that while the NPA applies to the salient individual, the DC does not. Thus two women, A and B, might have the following exchange about a potential suitor:

(61) A. I don't date stoners.

B. But Kyle doesn't smoke weed anymore!

A. Whatever, he's still a stoner. He lives in his parents' basement and eats pizza rolls all day.

Here, again, A felicitously doubles down on her use of the DC in the face of B's attempted correction.

⁷²See e.g. Is It Offensive to Call Someone a 'stoner'?:

"My issue is your use of 'stoner' as a rather pejorative term, whether derogatory or not, like the words 'faggot' or 'queer,'" one reader wrote me in an email, while accusing me of "painting all enjoyers with the abuser paint brush" and setting back the legalization movement."

⁷³I have been using single quotation marks to indicate that an expression is being mentioned *as* an expression, *viz.*, as a linguistic object of study. Here, 'stoner' is not being mentioned in this sense: it appears in a belief report about speakers who use the DC, and who presume that it has an actual extension. It is well-recognized, however, that the offensiveness of pejorative expressions can project out of attitude reports and other kinds of embeddings ????. To mitigate this effect as much as possible, I will place DCs that I do not strictly speaking mention, but wish not to seem to endorse, in double quotation marks.

⁷⁴See, e.g., this internet discussion. See also the Reddit discussion linked in footnote 3.

'libtard'/'liberal'

A portmanteau of 'liberal' and 'r*tard', 'libtard' is an offensive political pejorative used much the same way as the previous expressions.⁷⁵ Like 'mom jeans' and 'dad joke' (and less explicitly 'stoner'), 'libtard' bears a clear etymological connection to its NPA. And just as in those cases, competent users associate both it and 'liberal' with many (if not all) of the same stereotypes—e.g., *lacking common sense*, *having "bleeding hearts," being idealistic, being unreasonable*, and *enforcing political correctness*⁷⁶ These are, as far as competent users are concerned, the qualities of individuals whose political views are "r*tarded." Thus a supporter of Trump's "Zero Tolerance" immigration policy claimed:

(62) Libtards want us all to love and forgive the scumbags that commit these crimes.⁷⁷

Whatever else the relationship between 'liberal' and 'libtard' involves, it clearly involves *C*-overlap. But just as with the previous DC/NPA pairs, this relationship clearly also involves *pE*-overlap. Competent users believe that most liberals are "libtards," and vice versa.

Still, competent users frequently make exceptions. Those who say things like (12) and (42), for example, do not believe that every liberal is a "libtard":

(12) It's okay to be liberal, but it's not okay to be a libtard.⁷⁸

(42) I wasn't saying all liberals are libtards. They are not. But libtards do exist and they need to be called out on their bullshit.⁷⁹

At the same time, competent users seemingly do not believe that every "libtard" is a liberal. Suppose that A and B, for example, are talking about Joe online. B knows that Joe identifies politically as a libertarian, and is generally antagonistic to liberal views and policies. A, on the other hand, does not know much about Joe's political views. All A knows is that Joe recently posted on Facebook that he supports open borders. Now, suppose that they have the following exchange:

(63) A: Joe's a libtard.

B: But Joe's a libertarian, not a liberal!

A: Whatever, he's still a libtard. He thinks we should have open borders.⁸⁰

By uttering 'Whatever', A has effectively asserted (64):

⁷⁵Though 'libtard' is often considered a "juvenile" pejorative, it is far from mild. It inherits significant derogatory force from the highly offensive and ableist expression 'r*tard'.

⁷⁶The top definition on *UrbanDictionary.com* defines 'libtard' as follows:

noun. an individual, whose thinking process has been rendered impaired by political correctness and the failure to understand that people are responsible for their actions and the world does not owe lazy or stupid people a living.

For more, see the rest of the online entry.

⁷⁷Twitter, 18 March 2017.

⁷⁸Twitter, 2 September 2018.

⁷⁹Twitter, 20 Aug 2018.

⁸⁰Twitter, 19 Sep 2019.

(64) A: Joe's not a liberal, but he's a libtard.

And just as in the previous DC-corrections, A has felicitously doubled down. Competent users may find cases like Joe's *surprising*, though they will generally accept, on reflection, that they can obtain.

'c*nt'/'woman'

This is also intuitively true for 'c*nt'/'woman'. Speakers of American dialects, at least, will recognize this a clear DC/NPA pair.⁸¹ For example, consider the aggrieved male programmer who posts (65) and (66) in an online forum:

(65) Everyone can learn to program but the vast majority are [idiots] or some feminist affirmative action cunts who have no passion in engineering.⁸²

(66) Imagine lying about sexual assault to get me fired from my job, sick cunts.⁸³

With both uses of 'c*nt', the programmer intuitively targets women. And, intuitively, he targets these women *as women*. Yet with neither use of the DC does he mean to pick out the *general class* of women.⁸⁴ Indeed, that misogyny *makes exceptions* is one of Kate ?'s most important insights:

Misogynists can love their mothers—not to mention their sisters, daughters, wives, girlfriends, and secretaries. They need not hate women universally, or even very generally. They tend to hate women who are outspoken, among other things (52).

While competent users believe that many women are "c*nts," they do not generally believe that *all* women are "c*nts." They believe, rather, that only some women are—*viz.*, those who are, inter alia, *catty and mean, difficult and unreasonable*, and *outspoken*.

Moreover, as examples like (58) suggest, competent users sometimes also make exceptions in the other direction:

(58) Not all women are cunts. Term is reserved for only the deserving. Men can be cunts too. Just as they can be pussies.⁸⁵

Thus, as competent users of both expressions, A and B might have the following exchange:

(67) A: Ike's a cunt.

B: But Ike's a guy!

A: Whatever, he's still a cunt.

⁸¹In the United States, 'c*nt' is associated extremely closely with 'woman', and is considered one of the most offensive and taboo-ed expressions in the English language. This is not true, however, in other parts of the English-speaking world. In New Zealand, Australia, and parts of the United Kingdom, for example, it is used as a neutral or even positive expression. See, e.g., the Wikipedia entry.

⁸²Discussion forum, 21 Jun 2016.

⁸³Twitter, 13 November 2018.

⁸⁴These expressions, ? writes, "are typically applied to particular types of women, those whose actions, attitudes, or social stature defy misogynist norms and patriarchic entitlement. Mothers, sisters, wives, and daughters are the reason why most languages lack slurs targeting women wholesale" (n. 13).

⁸⁵Twitter, 11 Dec 2016.

In my experience, exchanges like this—whereby ‘c*nt’ is applied to a man not in a targeted, name-calling way, but as a descriptive predication—are more common in American gay communities. And here, again, *pE*- and *C*-overlap seem to capture why B’s confusion is understandable: competent users associate ‘c*nt’ *predominantly* with women, and generally do not expect to hear it applied to men. Still, they accept that some such applications are appropriate; and for this reason B’s correction is not conversation-stopping.

*‘d*ke’/‘lesbian’*

Precisely the same pattern, finally, seems intuitively to obtain for ‘d*ke’/‘lesbian’, the last DC/NPA in Group B. Competent users clearly presume substantial extensional overlap between the two expressions. Such speakers also associate the two expressions with many of the same stereotypes—such as *being a woman*, *being same-sex attracted*, *being butch/masculine*, and *being aggressive or outspoken*. But as in the previous cases, competent users are willing to make exceptions.

As far as such speakers are concerned, some lesbians (especially those who are feminine) are not “d*kes.” But competent users also believe that there can be “d*kes” who not lesbians—e.g., bisexual women. Thus A and B might have the following exchange:

- (68) A: Molly’s a dyke.
 B: But she’s bisexual, not a lesbian!
 A: Whatever, she’s still a dyke.⁸⁶

Once again, A’s doubling down is felicitous. Accepting that Molly is not a lesbian does not require A, as a competent user, to accept that she isn’t a “d*ke.”

By and large, DC/NPA pairs like these have escaped the attention of theorists working on slurs; and the extent that they haven’t, they have analyzed differently than pairs like ‘k*ke’/‘Jew.’ Notably, however, even the most paradigmatic slur/“neutral counterpart” pairs exhibit the same patterns. In particular, competent users of DCs like ‘f*ggot’, ‘ch*nk’ and ‘k*ke’ often have exchanges like (69)-(71):

- (69) A. You’re a faggot.
 B. But I’m bisexual, not gay!
 A. Whatever, you’re still a faggot.⁸⁷
- (70) A. All the nail techs here are chinks.
 B. But they’re Vietnamese, not Chinese!
 A. Whatever, they’re still chinks.⁸⁸
- (71) A. That lawyer’s a kike.
 B. But she’s not Jewish!
 A. Whatever, she’s still a kike.⁸⁹

⁸⁶Twitter, 1 Oct 2012.

⁸⁷Twitter, 10 Apr 2014.

⁸⁸Twitter, 5 May 2011, Twitter, 26 Oct 2009.

⁸⁹Twitter, 13 Sep 2013.

In each case, as in the previous DC-corrections, B's response is natural, but not conversation-stopping; and A's reassertion of the DC is felicitous, even as he in the same breath rejects the NPA. Clearly, competent users of such paradigmatic slurs associate many (if not all) of the same stereotypes with the relevant NPAs. And clearly, such users take there to be substantial *extensional* overlap between the relevant expressions. But just as with competent users of 'chick flick'/'romantic comedy', such speakers commonly make exceptions. And if an overlap thesis plausibly explains DC-corrections like (23), then seemingly it can also explain (69)-(71).

§6 Orthodoxy

It will be helpful, at this point, to take stock of the evidence. (1) For competent speakers of (American) English, there is a wide range of group or category expressions with recognizably pejorative uses (DCs). (2) Among these group or category expressions, many are such that competent speakers associate them intuitively with other, more "neutral" group or category expressions (NPAs). (3) Many such expressions seem to come in *pairs* (DC/NPA pairs), including paradigmatic slur/"neutral counterpart" pairs like 'k*ke'/'Jew', which pattern systematically in at least some common linguistic contexts. (4) And for many of these pairs — including chick flick/'romantic comedy', 'mom jeans'/'jeans worn by moms', 'stoner'/'cannabis user', 'libtard'/'liberal', 'c*nt'/'woman', and 'd*ke'/'lesbian' — this behavior is intuitively explained by an *overlap thesis* about the target linguistic relationship. What grounds the intuitive link between the DC and the NPA, in these pairs, is intuitively the presumption of (more or less complete) extensional overlap by speakers who use the expressions, grounded by the (more or less complete) conceptual overlap these users exhibit.

Together, (1)-(4) suggest a general picture of how DC/NPA pairs work. And on this general picture, DC/NPA pairs like 'k*ke'/'Jew', 'ch*nk'/'Chinese', and 'f*ggot/gay' work the same way that DC/NPA pairs like 'chick flick'/'romantic comedy' and 'liberal/libtard' do: speakers who use such expressions associate them with many of the same stereotypes, beliefs, and attitudes, and (so) presume them to apply to many of the same things. But as with other DC/NPA pairs, this doesn't tell us *how much* extensional overlap such speakers presume between the two expressions, or if *they're right*. Unbeknownst to slur users, the extensions they presume to pick out might be empty.

This general picture of DC/NPA pairs is simple, unified, and captures a wide range of the relevant data. It is surprising, then, that the resulting view of slurs and their so-called "neutral counterparts" is utterly absent from the existing literature.

Indeed, as far as I am aware, no one in the philosophical or linguistic literatures has even considered an "overlap view" for pairs like 'k*ke'/'Jew'. What's more, nearly all extant theories of slurs are *straightforwardly inconsistent* with the general picture of DC/NPA pairs just sketched. Despite their many points of disagreement, nearly all such theories assume that there is some *special* relationship for DCs like 'k*ke' and their NPAs that at once *unifies* pairs like 'k*ke'/'Jew', 'ch*nk'/'Chinese', and 'f*ggot/gay' but (in effect) *distinguishes them* from pairs like 'c*nt'/'woman', 'stoner'/'cannabis user', and 'chick flick'/'romantic comedy'. As we'll see shortly, this relationship is cashed out in different ways by different theorists. But in every case, this special relationship is assumed to give us the *meaning* of a given slur,

and (so) what is required for basic competence with it. Familiarity with stereotypes, it is now standardly held, is *not* required for competence with slurs—and certainly not “neutral counterparts.”⁹⁰

Among existing theories of slurs, all but one⁹¹ makes what I will call the *just-add-bad assumption* about the relationship between slurs and so-called “neutral counterpart” terms. Suppose that there is a well-defined (call it \mathbb{S}/\mathbb{N}) set of all and only slur/neutral counterpart pairs. The *just-add-bad* assumption (hereafter JAB) says that for any pair of expressions s/n , if $s/n \in \mathbb{S}/\mathbb{N}$, then there obtains between s and n a semantic relationship r such that:

- (a) the truth-conditional meaning of n can be used to analyze, conjunctively or identificationally, the proposition (if there’s only one) or propositions (if there’s more than one) expressed by assertoric sentences containing s , and
- (b) recognition of r , together with knowledge of the truth-conditional meaning of n and knowledge that s is pejorative, is sufficient for competence with s .

Most theories assume, in other words, that there is a *recipe* for computing the truth-conditional meanings of slurs from those of their so-called neutral counterparts; and that this recipe, once we “add the bad”, makes competence with slurs easy to come by.

JAB is clearest in what Adam ? calls *coreferentialism*, according to which slurs and “neutral counterparts” are *truth-conditionally equivalent*.⁹² This is the thesis that expressions like ‘k*ke’ and ‘Jew’, as a matter of general semantic fact, refer to the exact same groups, and interact with truth-functional operators like ‘and’ in the exact same way.⁹³ Thus, coreferentialism says that if we know the extension of ‘Jew’, then we know the extension of ‘k*ke’. And if that’s all there is to meaning, then (given we recognize that ‘k*ke’ is pejorative) we are competent with the slur.⁹⁴ The coreferentialist “recipe” is very straightforward.

Coreferentialist views are what I will call *Identificational Views* of slurs. According to Identificational Views, the truth-conditional component of slurs, or the part that determines extension, is given *identificationally* by the extension of the corresponding neutral counterpart term. Slurs either make the *same* propositional contributions as their neutral

⁹⁰Some early theories (most notably ?) built stereotypes into the meanings of slurs; however, since ?’s influential critique, stereotype views have fallen out of favor.

⁹¹Neufeld:2019vu’s essentialist prototype view of slurs is a lone exception.

⁹²? rejects coreferentialism on much the same grounds as I ultimately want to (*viz.*, on the basis of sentences like (52), “He isn’t a k*ke like 99% of the Jews.” Still, even Croom still accepts a version of JAB, as the “conceptual anchors” of slurs on his prototype semantics are given by the relevant “neutral counterparts” (35).

⁹³Cf. ?, who calls this Neutral Counterpart Theory. All of the views cited in footnote 2 are coreferentialist. This includes ?, though he denies that sentences with slurs are truth-apt.

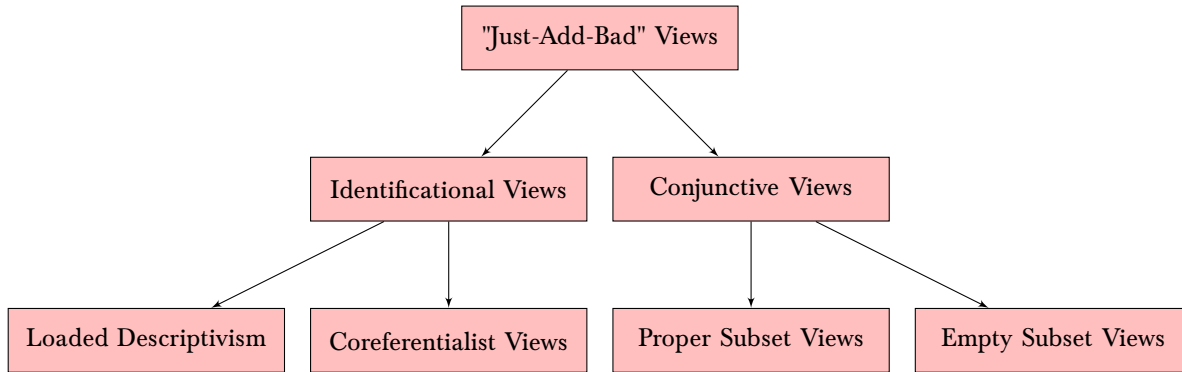
Cf. ?who, in their response to ?, break coreferentialism into two claims, which I will rename (C1) and (C2):

(C1) For every slur s , there is a neutral counterpart n with the same extension.

(C2) The contribution made by a slur s to the truth-conditional content of a sentence S in which it occurs is the same as the contribution made by its neutral counterpart n to $S[n/s]$, where $S[n/s]$ is the sentence that results from S by substituting n for s in S one or more times (265-6).

As they observe, (C1) is required if coreferentialist theories are to apply to all slurs.

⁹⁴There may be other features of ‘kike’ that distinguish its meaning from that of ‘Jew’—like an expressive dimension, a taboo, or a pragmatically generated message or signal; but at their truth-conditional “semantic core”, the expressions are precisely coextensive.



counterparts (*à la* coreferentialism); or they make *more than one* propositional contribution, the contents of which are fixed by the relevant neutral counterparts. Identificational Views thus include Kent ?’s *Loaded Descriptivism*, according to which the meaning of a slur, *s*, involves the meaning of its neutral counterpart, *n*, plus a “side comment” along the lines of “[*n*]*s* are contemptible in virtue of being [*n*]” (p. 64).⁹⁵ (This is essentially the same “recipe” as coreferentialism, with “the bad” added in via a second proposition).

Most extant theories of slurs are coreferentialist, and I will be focusing on them in what follows.⁹⁶ But there are other, non-identificational views that nevertheless accept JAB. These are what I will call *Conjunctive Views*.

Conjunctive Views analyze the truth-conditional meanings of slurs in terms of *being n* (*viz.*, being in the extension of the relevant counterpart term) *and* something else. Examples include the view, proposed by ??, that slurs predicate the property *ought to be the target of negative moral evaluation because of being a member of G*, where *G* is the relevant “neutral counterpart” group; as well as the view that the extensions of slurs are *proper subsets* of the extensions of their neutral counterparts.⁹⁷ Thus slurs are not, on conjunctive views, necessarily truth-conditionally equivalent with their neutral counterparts, and indeed may not have any extension at all.⁹⁸

The differences among and within these families of views are numerous and complex. For my purposes here, however, the views in each have one important thing in common: they are facially false when extended to the DC/NPA pairs discussed in the previous two

⁹⁵Bach also explicitly leaves it open whether slur ascriptions are “true of whomever [relevant] neutral counterpart is true of” (61).

⁹⁶These include: (a) *hybrid expressivist views*, according to which the meanings of slurs involve the descriptive meanings of their NPA plus a negative attitude ??? or gestural content ?; (b) *implicature views*, according to which slurs conventionally imply something negative that their NPAs do not (e.g., ???); (c) *prohibition views*, according to which the only difference between slurs and their NPAs is taboo (e.g., ?, ?); and (d) *Contrastive Choice views*, according to which uses of slurs and NPAs are truth-conditionally equivalent, and uses of slurs where NPAs are available are “pointed choices” signaling endorsement of associated attitudes and ideologies (e.g., ??).

⁹⁷This view is not uncommonly voiced by non-philosophers, and is sketched by ?.

⁹⁸Since it is true of no one that they *ought to be the target of negative moral evaluation because of being a member of G*, where *G* is a relevant “neutral counterpart” group, ?? hold that slurs have null extension. Similarly, ? previously held that slurs predicate the complex property “ought to be subject to discriminatory practices [XYZ] because of having negative properties [XYZ], all because of being [*n*],” and (since no one satisfies that property) are actually empty.

sections.

Consider again some of the bidirectional divergence data presented in §4:

- (3) The new *Ghostbusters* is chick flick, but it isn't a romantic comedy.
- (4) *Silver Linings Playbook* is a romantic comedy, but it isn't a chick flick.
- (32) Being a dad with "actually funny jokes and not just dad jokes" might be the best feedback you can get from your kid.
- (33) I don't understand why Cory Booker tells so many Dad jokes when he's not actually a Dad
- (26) Grandma is not a boomer, she's cool.
- (27) Mayo Pete is a boomer at 37. Does this help everyone understand that boomer isn't just an age thing?
- (46) Just because someone is a cannabis user doesn't mean they're a stoner.
- (47) Just because someone isn't a cannabis user doesn't mean they're not a stoner.
- (54) There are gays, and then there are faggots.
- (55) *said about an executive ban on transgender military personnel:*
Trump won't allow faggots in the military. [But] I say if they can do the job of a soldier, let 'em serve.
- (56) I always wanted a lesbian friend, lesbian not a dyke.
- (57) Not all dykes are lesbians. I got a cousin who's a dyke but she has a husband.
- (58) Not all women are cunts. Term is reserved for only the deserving. Men can be cunts too. Just as they can be pussies.

The range of these examples, together with the systematic behavior observed in §5, gives us *pro tanto* reason, at least, to treat the relevant DC/NPA relationship the same way in each case. If we take this *pro tanto* reason seriously, however, then is clear that neither coreferentialist nor non-coreferentialist JAB views will be tenable as a general theory. It is obviously false that 'chick flick' and 'romantic comedy' are coextensive; obviously false that 'chick flick' means "contemptible in virtue of being a romantic comedy"; and obviously false that "chick flicks" constitute a proper subset of "romantic comedies." And the same theses are similarly obviously false about 'dad joke'/'joke told by dads', 'mom jeans'/'jeans worn by moms', 'stoner'/'cannabis user', 'libtard'/'liberal', 'c*nt'/'woman', and 'd*ke'/'lesbian', respectively, among other DC/NPA pairs listed in §2 like 'boomer'/'(American) person born 1946-1964' and 'Bernie Bro'/'Bernie Sanders supporter'.

Of course, it would be far too hasty to conclude from this that JAB views are false. What it does show, however, is that there is a wide range of data that JAB views must be able to accommodate in order to be plausibly true. And I think there is reason for skepticism that such accommodations can be made in non-*ad hoc* ways. In the next section, I'll present a dilemma for coreferentialists views in particular, though the same considerations apply to non-coreferentialist views.

§7 Dilemma

As I see it, coreferentialists have two options for explaining the apparent systematicity across DC/NPA pairs. First, they can explain it in terms of a *distinct* linguistic mechanism,

which operates systematically in the relevant problem cases. In my experience, the mechanism most commonly reached for involves metaphorical or otherwise non-literal uses of DCs, and that is what I will be focusing on. The second option is to postulate a primitive semantic difference between the problem cases and the non-problem cases regarding the relevant DC/NPA relationship. As I hope to show, however, both of these options are dubiously *ad hoc*, and moreover unnecessary: a basic overlap thesis for DC/NPA pairs can get us the same benefits with none of the costs.

Option 1: Appeal to an independent systematic mechanism

To begin, recall (23), which I called a *DC-correction*:

- (23) A: The new *Ghostbusters* is a chick flick.
 B: But it's an action film, not a romantic comedy!
 A: Whatever, it's still a chick flick.

Now, compare (23) with (72), which I will call an *NPA-correction*:

- (72) A: The new *Ghostbusters* is a romantic comedy.
 B: But it's an action film, not a romantic comedy!
 A: # Whatever, it's still a romantic comedy.

(72) is just like (23), except the instances of 'chick flick' have been replaced with 'romantic comedy'. But unlike (23), this exchange is defective. Whereas before A could *feliculously* double-down in the face of B's correction, his rejoinder here is bizarrely uncooperative. A natural explanation for this crash is that, in accepting B's correction ('Whatever'), A is intuitively conceding that the new *Ghostbusters* is not a romantic comedy. And A cannot coherently concede this while maintaining his original claim—*viz*, that *Ghostbusters* is a romantic comedy. Confronted with this level of uncooperativeness, it would be eminently forgivable for B to abandon the conversation completely.

But now, suppose that we were coreferentialists about 'chick flick'/'romantic comedy' and thought that 'chick flick' and 'romantic comedy' were coextensive in their literal uses. On this basis alone, and all other things being equal, we might expect A's rejoinder in (23) to feel as infelicitous and conversation-stopping as his rejoinder in (72). So given that it *doesn't*, how might we explain the felt asymmetry?

Most likely, we would hypothesize that A is using 'chick flick' in (23) in a *nonliteral* or *nonbasic* way, such that it refers to a more *expansive* class than 'romantic comedy' in (72). Perhaps A intends such an expanded meaning throughout the entire conversation with B, or perhaps he *shifts* to a non-literal meaning midway through. After all, we can easily imagine exchanges working this way, such as (73):

- (73) A: Jack's a girl.
 B: But Jack's a boy!
 A: Whatever, he's still a girl. He cries all the time and can't take a joke.⁹⁹

⁹⁹This is a modified version of an example from ?.

In (73), A clearly intends a nonliteral, expanded use of ‘girl’ from the outset. It is no surprise, then, that his doubling-down is felicitous despite his concession Jack is a boy.

This kind of move has been informally suggested to me for DC- and NPA-corrections like (70) and (74), and has been proposed in print by ? for sentences like (54) and (55):¹⁰⁰

- (70) A. All the nail techs here are chinks.
 B. But they’re Vietnamese, not Chinese!
 A. Whatever, they’re still chinks.

- (74) A. All the nail techs here are Chinese.
 B. But they’re Vietnamese, not Chinese!
 A. # Whatever, they’re still Chinese.

(54) There are gays, and then there are faggots.

(55) *said about an executive ban on transgender military personnel:*

Trump won’t allow faggots in the military. [But] I say if they can do the job of a soldier, let ‘em serve.

But is it really plausible that the same thing is going on in these cases, much less (23)? My intuition, very strongly, is no. As competent speakers, we *know* that A intends a nonliteral meaning of ‘girl’ in (73). Moreover, we have a strong sense of the *way* that B is misunderstanding A when he attempts to correct him—*viz.*, by taking as literal an utterance which is clearly supposed to be figurative. But we, theorists who (presumably) do not ourselves use slurs, should not be deceived by the structural similarities into thinking that the same is plausible in cases like (70). A main reason I have been using such a wide range of examples in is because I want to include expressions that readers actually use. And users of ‘chick flick’ know that a figurative or otherwise “nonbasic” story is not plausible for (23), just as users of ‘boomer’ know that it is not plausible for (75):

- (75) A. Mayo Pete is a boomer.
 B. But he was born in 1982!
 A. Whatever, he’s still a boomer.¹⁰¹

The core practice of most actual slur users, I submit, is not targeted name-calling; it is grotesquely mundane predication in conversations with one another. It is a practice of distinguishing “model minorities” from “the riff raff”; of deciding who and who should not “count.” That so many theorists think otherwise is, I suspect, the result of their taking a very extreme kind of user as their paradigm. These are the avowed white supremacists, the proud and evangelizing bigots, whose hatred for particular groups is not only profound, but *directed* and *absolute*. In their hands, DCs like ‘n*gger’ are weapons reserved for *all and only* the members of the hated groups—groups which we would otherwise pick out with NPAs. But most racists are not so unyielding in their contempt. Indeed, the willingness of

¹⁰⁰Jeshion calls these “G-contracting” and “G-extending” uses of slurs, respectively, where “G” is the target group picked out by a slur and its neutral counterpart in literal uses (pp. 251-253).

¹⁰¹Twitter, 20 Nov 2019.

everyday racism to “make exceptions” is what makes it so insidious—and, in my view, what makes the coreferentialist thesis so absurd.

Having grown up around slur users, I *think* I know what they mean when they say things like (54) and (55):

(54) There are gays, and then there are faggots.

(55) *said about an executive ban on transgender military personnel:*

Trump won’t allow faggots in the military. [But] I say if they can do the job of a soldier, let ‘em serve.

That these speakers—like A in (73) with ‘girl’—are *knowingly* using DCs in extended or contracted ways strikes me as wildly implausible. In any case, though, I *know* that this hypothesis is implausible about (23). When I, as a competent user, reflect on A’s use of ‘chick flick’, I do not at all feel the need to reach for a nonliteral interpretation to make sense of what is said. On the contrary, it *really feels* like A’s meaning in (23) *just is* the meaning of ‘chick flick’ (whatever else we want to say about it). And the same seems true, on reflection, about the DCs used in (63), (67), and (68):

(63) A: Joe’s a libtard.

B. But Joe’s a libertarian, not a liberal!

A: Whatever, he’s still a libtard. He thinks we should have open borders.

(67) A: Ike’s a cunt.

B: But Ike’s a guy!

A: Whatever, he’s still a cunt.

(68) A: Molly’s a dyke.

B: But she’s bisexual, not a lesbian!

A: Whatever, she’s still a dyke.

So even if an appeal to nonliterality makes sense in the case of ‘girl’ in (73), it does not seem like the right thing to say about seemingly structurally parallel uses of many other DCs. As a coreferentialism-friendly way of explaining the systematicity data, it does not seem like it’s going to work.

Option 2: Postulate a primitive semantic distinction

Another tack the coreferentialist might take is to reject the systematicity data as *prima facie* reason to pursue a general theory of the DC/NPA relationship, and to postulate instead a primitive semantic distinction between DC/NPA pairs of which coreferentialism is true, and DC/NPA pairs of which it is false. That is, to postulate a subclass of DCs such that, as a brute fact about the sorts of expressions they are, they bear a *characteristically* coreferentialist relationship to their NPAs.

In spirit, this is actually not a novel proposal. A similar strategy has been employed by coreferentialists to defang objections from ? and ?, who have argued that coreferentialism, in failing to account for derogatory expressions like ‘slanty-eyed’, ‘curry muncher’ and ‘Jewish American Princess’ (DiFranco) and ‘slut’, ‘sissy’, and ‘b*tch’ (Ashwell), is simply too narrow to be a plausible theory of “slurs.” These objections have placed little pressure on coreferentialists, however, as they have felt justified in denying that such putative coun-

terexamples belong to the class of expressions they are concerned to analyze.¹⁰² After all, there is nothing objectionable about a narrow theory of expressions, if those expressions in fact make up a distinct semantic class.

Do paradigmatic slurs ‘k*ke’ and ‘ch*nk’ make up a distinct semantic class of DCs? Maybe.

It is possible, for example, that paradigmatic slurs factor differently than other DCs into attitude reports (where they might seem to project differently), or that they are lexicalized in distinctive ways.¹⁰³ And while I am not inclined myself to think such things, I have no aim of arguing otherwise, here. What I am interested in, to be clear, is not whether paradigmatic slurs are semantically distinct other DCs, but whether the *relationship* that they bear to their intuitive “neutral counterparts” is *essentially different*, as a brute semantic fact, from the relationship that other DCs bear to *their* intuitive NPAs. And it is this latter thesis that I think the coreferentialist needs if she is to retain her assumption, undercut by the systematicity data in §4 and §5, that all expressions like ‘k*ke’, *in virtue of being the sort of expressions they are*, corefer with expressions like ‘Jew’.

The best thing one can say for this thesis, though, is that it is not provably false. By all other accounts, it seems transparently *ad hoc*. There is no evidence for it that cannot be accommodated by an overlap view; and there is significant evidence against it. In the absence of obvious defeaters, the systematicity data in §4 and §5, together with the robust asymmetry between DC- and NPA-corrections like (23) and (72), favor a unified theory. On top of that, though, some of the DCs that coreferentialists already wish to capture are *a lot like* DCs to which their views obviously do not extend. To rescue her theory, then, the coreferentialist would need to posit some intuitively very implausible brute distinctions.

Consider, for example, ‘commie’ and ‘libtard’, which are both pejorative expressions associated with leftist political ideologies. Some coreferentialists (perhaps most notably ?) already assume that ‘commie’, an expression akin to ‘k*ke’ and ‘f*ggot’, is (*ipso facto*) co-extensive with the more neutral ‘communist’. The analog thesis about ‘libtard’/‘liberal’, however, is patently false. So in order to rule *in* ‘commie’/‘communist’ but rule *out* ‘libtard’/‘liberal’, the coreferentialist must postulate some essential difference between the two DC/NPA pairs such that (as it were) her view is true of one of them, but not true of the other.

But why should we believe this? Why should we believe, in particular, that ‘commie’/‘communist’ is essentially more like ‘k*ke’/‘Jew’ than ‘libtard’/‘liberal’? Such a semantic partitioning is not, of course, impossible. Language is incredibly complex, and unity is not the only theoretical value. But given the intuitive similarities between these DC/NPA pairs, it would be very surprising if the *reason* that an assumption of coreferentialism was licensed for one DC/NPA pair but not the other was that ‘commie’ and ‘libtard’ belonged to distinct semantic classes. Such a partitioning, at least, cannot be taken for granted. More importantly, it is not a bullet we have to bite.

The crucial thing about coreferentialism (and what I have called JAB views more

¹⁰²?, for example, argue in their response to DiFranco that slurring uses of expressions like ‘curry muncher’ have non-compositionally determined idiomatic meanings. ? claims that ‘bitch’ and ‘slut’ are gendered *personal* pejoratives, like ‘bastard’, that target individuals on the basis of personal traits and actions rather than group membership (54-55). Cf. Jeshion (2018), cited in footnote 84, above.

¹⁰³For skepticism about the semantic distinctiveness of paradigmatic slurs, see ?, especially pp. 240-244.

broadly) is that it purports to be a *general* theory of a certain *class* of terms. It *takes it for granted* that, for *any* DC relevantly like ‘k*ke’, that DC *will* bear the same coreferential relationship to an NPA that (*ex hypothesi*) ‘k*ke’ bears to ‘Jew’. Thus coreferentialism does not merely claim that *some* DCs are NPAs are, *as an accidental fact*, coextensive. An overlap thesis, however, *can* claim this for some DC/NPA pairs, should we want it to. After all, all an overlap thesis says is that the basic DC/NPA relationship involves both pE-overlap and C-overlap—and *complete* overlap is, in either case, of course a possibility.¹⁰⁴

§8 Diagnosis and Conclusion

This leaves us with the question: if coreferentialism, and JAB more generally, is so implausible as a general thesis, why has it been so popular?

For very good reasons, theorists interested in the semantics and pragmatics of slurs have generally aimed to explain their derogatory force. Indeed, if we are going to theorize about expressions which have observably harmful effects, we ought to do so in a way that teaches us something that we cannot learn from other, less harmful expressions. It is quite right, then, that theorists have generally aimed to explain slurs’ characteristic power to derogate and offend. But considerations of offensiveness tend to favor extreme cases; and by fixating on slurs’ derogatory force, theorists have imposed an artificial constraint on what counts as “relevant.” Thus it is only the very *worst* slurs, as used by the very *worst* bigots, that have been presumed central cases for theorizing about slurs.¹⁰⁵

These are what ? aptly calls “weaponized” uses of slurs, or cases like (69):

(69) *Shouted at a gay couple holding hands:*
You’re going to hell, faggots!

Such cases are defined principally by the individual(s) targeted, who have three main features: (i) they belong to the relevant NPA group; (ii) they are *believed* by the DC user to belong to the relevant NPA group; and (iii) they are targeted by the DC user *because* they belong to the relevant NPA group. Orthodox views, and in particular coreferentialist ones, are well-suited to explain (iii).

But this is a naïve view; and given the role of slurs in *perpetuating* exclusion and oppression, it would be surprising if it were right. As I have tried to emphasize, most ordinary slur use comes from ordinary racists, and ordinary racism “make exceptions.” Indeed, this is precisely Kate Manne’s critique of traditional conceptions of misogyny as “hatred of women.” There is an important reason that, according to the “logic” of misogyny, not all women are “c*nts”; and that some men (especially gay men) are. By distinguishing the

¹⁰⁴For ways to fill out the details of an “overlap view” see footnote 36, above.

¹⁰⁵? makes a similar observation, writing:

Writers focus almost entirely on what slurs convey about their targets and the insult or offense they give, not on what they have to say about the groups that coin and use them, though those group-identifying or group-affiliating uses are more prevalent, more universal, and arguably prior to their uses as terms of direct abuse. The motivations of the people who use slurs are pretty much discharged by describing the prototypical speaker as “the racist.” (241)

“good” women from the “bad” ones, misogyny insulates itself from charges of absolute sexism while (re)tightening its ideological grip.

The everyday racism and bigotry of slur users works in importantly analogous ways. This bigotry is essentially exception-making, relying for its survival on the possibility of admitting “good ones.” And we need not be slur users ourselves to see this! As I have tried to show, less potent (and more familiar) ideologies, such as those underpinning the use of ‘dad joke’, ‘stoner’, ‘chick flick’ and ‘boomer’, work intuitively the same way—*viz.*, by positing two closely related but (in most cases) ultimately nonidentical categories, characterized by closely related but (in most cases) ultimately nonidentical stereotypes. An overlap thesis about DC/NPA pairs can capture this general phenomenon without imposing the (apparently) arbitrary semantic constraints of the “just-add-bad” assumption. But rejecting the idea that there is any general, straightforward “recipe” for computing the meanings of paradigmatic slurs from their so-called “neutral counterparts” does not merely lead us to better semantic and pragmatic theories. More importantly, it positions us to better understand, identify, and confront the insidious mechanisms of ordinary bigotry.