

BUSTING THE GHOST OF NEUTRAL COUNTERPARTS

JEN FOSTER

University of Southern California

Slurs have been standardly assumed to bear a very direct, very distinctive semantic relationship to what philosophers have called “neutral counterpart” terms. I argue that this is mistaken: the general relationship between paradigmatic slurs and their “neutral counterparts” should be assumed to be the same one that obtains between ‘chick flick’ and ‘romantic comedy’, as well a huge number of other more prosaic pairs of derogatory and “less derogatory” expressions. The most plausible general relationship between these latter expressions — and thus, I argue, between paradigmatic slurs and “neutral counterpart” terms — is one of overlap in presumed extension, grounded in overlap in associated stereotypes. The resulting framework has the advantages of being simple, unified, and, unlike its orthodox rivals, neatly accommodating of a much wider range of data than has previously been considered. More importantly, it positions us to better understand, identify, and confront the insidious mechanisms of ordinary bigotry.

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 a 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*:

1. The [trailer for the 2016 Ghostbusters](#) was “disliked” on YouTube over 600,000 times before the film was even released, becoming the site’s “[most disliked movie trailer ever](#)” (Huddleston 2016).

2. See, e.g., “[Ghostbusters Director Says ‘Chick Flick’ Is a ‘Derogatory’ Term](#)” (Lasher 2016).

Contact: Jen Foster <jennifnh@usc.edu>

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.”

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 that ‘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 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 called the new *Ghostbusters* a chick flick, presumably no one called it a romantic comedy. Indeed, though competent users believe that ‘chick flick’ and ‘romantic comedy’ apply typically to the same things, they accept that cases like (1) and (2) can, at least in principle, obtain:⁴

- (1) The new *Ghostbusters* is chick flick, but it isn’t a romantic comedy.
- (2) *Silver Linings Playbook* is a romantic comedy, but it isn’t a chick flick.⁵

My purpose in this paper is not (just) to theorize about the terms ‘chick flick’ and ‘romantic comedy’. Rather, ‘chick flick’ and ‘romantic comedy’ will be my way into investigating a much broader class of expressions, the most important of which are slurs.

3. Throughout 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 (for example, because they grew up around people who use it), but do not themselves use it, or use it anymore.

4. That does not necessarily mean that speakers always *agree* about which things belong to which category. For examples of speakers debating the boundaries of the two terms, see [this Reddit discussion](#) dedicated to the difference between ‘chick flick’ and ‘romantic comedy’.

5. In response to the question “Is *Silver Linings Playbook* really a romantic comedy?”, one commenter in the discussion just mentioned (see footnote 4, above) writes:

That’s accurate. *Silver Linings Playbook* follows the formula pretty closely. It even ends with a dance number. It still has a comedic tone, just darker than the kind you usually get with a romance movie. But that’s in order to fit what’s supposed to fit the psychology of the leads, which gives them some more plausible conflicts and obstacles than you usually get in this kind of movie. It can be seen as an intelligent filmmaker trying to revise a genre by working within it. What kind of movie did you think it was?

Along similar lines, a number of readers have independently suggested (n1):

- (n1) *Die Hard* is a romantic comedy, but it isn’t a chick flick.

Slurs,⁶ I will suggest, bear the same relationship to so-called “neutral counterpart” terms as ‘chick flick’ does to ‘romantic comedy’. In particular, we should at least start by assuming that, semantically, sentences like (3) and (4) are akin to (1) and (2):

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

My purpose in this paper is use this starting assumption to dismantle what we’ll see is a thoroughly orthodox idea about pairs (1)/(2) and (3)/(4) in the philosophical slurs literature. I will then propose that we replace this idea with what I’ll call an *overlap thesis* about pairs of expressions like ‘chick flick’/‘romantic comedy’ and ‘k*ke’/‘Jew’. The resulting framework has the advantages of being simple, unified, and, unlike its orthodox rivals, neatly accommodating of a much wider range of data than has previously been considered.

Most importantly, however, the overlap thesis captures something about slurs and the people who use them which has been woefully obscured by the existing literature: that everyday bigotry makes exceptions.

The vast majority of the philosophical slurs literature, I submit, has taken the wrong cases as central—viz., cases like (5):

- (5) [Shouted at a gay couple holding hands]
 You’re going to hell, faggots!

These are what Robin Jeshion (2013b), aptly, calls “weaponized” uses of slurs—they are attacks. They are (among other things, which I say more about later) characteristically second-personal, extremely socially aggressive, intensely direct uses of slurs. And it makes sense for philosophers to care about such uses—they have extraordinary potential for harm.

But explanations of offensiveness and harm are only one desideratum of a theory of slurs; it should also elucidate the beliefs and attitudes of the people who use them. Most ordinary slur users, though, are not virulently absolutist in

6. Throughout this paper, I have attempted as much as possible to minimize how frequently such expressions are mentioned; and, where mentioning them is necessary, to blunt the force they inevitably carry. The worst and most offensive ones (a subjective distinction which, admittedly, is arbitrary) are transcribed with an asterisk (*) wherever they occur, *except* where they appear as part of in-set example sentences. Nearly all such in-set examples come verbatim from online sources; footnotes have been provided containing direct (at the time of publication) or otherwise archival directs to the original conversational contexts. Other examples have been edited for length or clarity, or else to minimise extraneous occurrences of slurs. Where such sentences have been adapted from online source data, this is indicated in the corresponding footnote.

7. Adapted from: Twitter, 22 Mar 2018. [[direct](#)] [[archive](#)]

their bigotry—and most ordinary slur use is not like (5). Theorists, in focusing on offensiveness, have mistakenly centered the very worst slurs, as used by the very worst bigots. Consequently, philosophical orthodoxy has failed to capture not only the speech, but the thought, of ordinary slur users. Such speakers are in the business—as in (3) and (4)—of asserting, arguing about, and admitting “exceptions to the rule”; this is not a peripheral feature of ordinary bigotry, but its core.

2. Setup

2.1. Fixing Ideas

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 Nunberg (2018), 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 unwarrantedly, 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” assumptions about slurs—viz., that expressions like ‘k*ke’ and ‘ch*nk’, in virtue of being slurs, bear a unique, semantically-given relationship to non-slur expressions like ‘Jew(ish)’ and ‘Chinese’.

8. Nunberg 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” (2018: 239). This view is explicitly suggested by Davis and McCready (2020).

9. In this vein Nunberg (2018) 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. (2018: 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 (McKinnon 2018; Allen et al. 2018).

10. For a nice discussion, see Diaz-Legaspe (2019).

The assumption I have in mind, here, is so foundational in the philosophical literature, and lurks so far in the discursive background, that pinning it down precisely is actually a little tricky; I will take up that task in §9. For now, just to help fix ideas, we can start with a paradigmatic example of a view that is committed to this assumption: Timothy Williamson's (2009) conventional implicature account of slurs and so-called "neutral counterparts". (We will see how to generalize from Williamson's view later.)

Williamson proposed that expressions like 'Jew' are, in a very literal sense, the "neutral" "counterparts" of derogatory expressions like 'k*ke': they have exactly the same truth-conditional content, just not the same conventional evaluative upshot. 'K*ke' and 'Jew' refer to exactly the same group of people, but differ in what (else) they imply or communicate about that group. Thus, following Dummett's (1973) discussion of the pejorative 'Boche', Williamson writes:

[T]o assert 'Lessing was a Boche' would be to imply that Germans are cruel, and I do not want to imply that, because the implication is both false and abusive. Since the false implication that Germans are cruel does not falsify 'Lessing was a Boche,' it is not a logical consequence of 'Lessing was a Boche.' Rather, in Grice's terminology, 'Lessing was a Boche' has the conventional implicature that Germans are cruel, in much the same way that 'Helen is polite but honest' has the conventional implicature that there is a contrast between Helen's being polite and her being honest. Just as 'Lessing was a Boche' and 'Lessing was a German' differ in conventional implicatures while being truth-conditionally equivalent, so too 'Helen is polite but honest' and 'Helen is polite and honest' differ in conventional implicatures while being truth-conditionally equivalent. (2009: 149–50, emphasis in the original)

In other words, "neutral counterparts" are perfect semantic proxies for the (truth-conditional) meanings of slurs: it is enough to know, for example, which group 'German' picks out to know which group 'Boche' does (and vice versa), 'Chinese', 'ch*nk', and, crucially, so on and so forth for all other relevant pairs of expressions. If Williamson's conventional implicature view is correct, it is a general semantic fact about sentences like (6) and (7) that they stand and fall together:

- (6) Isaiah is a Jew.
- (7) Isaiah is a kike.

Williamson takes it explicitly for granted that this is something "competent" speakers know, or "are in a position to know" (2009: 149). Many philosophers, most of whom (presumably) do not themselves use slurs, or associate closely with those who do, have followed his lead in accepting this "fact" as given. This has been a mistake—one that has led to serious distortions about what counts as "basic competence" with slurs, as well as what that competence requires.

We become much better positioned to see this, I suggest, once we broaden our focus to include expressions which we ourselves (or people whom we encounter in our day-to-day speech communities) do use. As I will show, many such expressions pattern systematically with the kinds of paradigmatic slurs philosophers have tended to focus on, and so (I submit) should be accommodated by any adequate semantic theory. Still, we might hesitate to call many of these additional expressions “slurs”; or indeed, for the sorts of definitional considerations already mentioned, even positively deny that they are ones. I wish to remain maximally neutral on this front. I will thus be setting aside the word ‘slur’ for the majority of what follows. I will talk instead of what I’ll call derogatory classifiers.

2.2. *Derogatory Classifiers and Non-Pejorative Associates*

As I intend the term, “derogatory classifier” (hereafter “DC”) covers a much wider range of (more or less) derogatory expressions than has typically been considered in philosophical work on slurs. In addition to all of the expressions which theorists have tended to focus on—paradigmatic slurs such as ‘k*ke’, ‘ch*nk’, ‘n*gger’, ‘f*ggot’—DCs also include many expressions which theorists have tended only to mention in passing, and many more which they have never considered at all. Here are some examples:

‘alchie’	‘d*ke’	‘junkie’	‘scab’
‘anti-vaxxer’	‘dad bod’	‘Karen’	‘shrink’
‘backwater’	‘dad joke’	‘libtard’	‘slut’
‘Bernie bro’	‘dive bar’	‘man cave’	‘soy boy’
‘Bible banger’	‘feminazi’	‘McChurch’	‘sp*c’
‘bimbo’	‘fleabag motel’	‘McMansion’	‘tankie’
‘boomer’	‘flyover state’	‘mom jeans’	‘tech bro’
‘bootlicker’	‘frat bro’	‘neckbeard’	‘tourist trap’
‘breeder’	‘gamer’	‘nerd’	‘towelhead’
‘breeder bar’	‘gangbanger’	‘parasite’	‘townie’
‘cape(shit) movie’	‘gas guzzler’	‘pig’	‘trailer trash’
‘chav’	‘geek’	‘pill mill’	‘tr*nny’
‘cheesehead’	‘ginger (kid)’	‘pillhead’	‘truscum’
‘chick flick’	‘goy’	‘plebe’/‘pleb’	‘traphouse’
‘coastie’	‘gringo’	‘poof’	‘treehugger’
‘commie’	‘hillbilly’	‘poser’	‘Trumper’
‘cracker’	‘hobo’	‘rag’	‘tourist trap’
‘cripple’/‘crip’	‘hole in the wall’	‘rainbow corp’	‘wetback’
‘cuck’	‘horse girl’	‘redneck’	‘woke’
‘c*nt’	‘jarhead’	‘r*tard’	‘yankee’
‘curry muncher’	‘jock’	‘rug muncher’	‘yuppie’

I do not expect the reader to recognize every expression on this list, or to agree with me that every expression on it deserves to be called “derogatory.” Hopefully, though, enough of the expressions are sufficiently familiar for the reader to detect a pattern. It has been my experience that, when presented with a handful expressions on the foregoing list, ordinary, competent (American) English speakers can spontaneously, effortlessly supply others to add to it. This, I take it, is strong *prima facie* evidence that DCs form some kind of unified linguistic class, whose members include at least many of the above expressions. (If you don’t think an expression from my list belongs to this class, please just put it to the side and interpret everything I say as going for the rest of the expressions, as well as for other members of the class.)

Importantly, though, in assuming that DCs form some kind of unified class, we need not assume anything about what makes that class unified. In the first place, as I’ve already emphasized, I do not wish 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; it is an advantage of DC-talk that it does not (or at least need not) carry the same weight.

Likewise, and perhaps more to the point given my purposes here, I do not wish to beg any questions about DCs’ semantic status. Broadly, DCs are nominalized, variably pejorative¹¹ (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” (Bach 2018: 60). But in saying that they are “pejorative” or “derogative”, I mean only that DCs have established pejorative or derogative uses in at least some speech communities.¹² Whether those uses are explained by

11. Derogatory variation—i.e., that pejorative expressions can vary in their derogatory force or intensity—has been long-observed in the philosophical literature on slurs (e.g., Hom 2018: 426; Jeshion 2013b: 233; Nunberg 2018: 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.

12. Theorists often distinguish between “group slurs,” like ‘k*ke’ and ‘commie’, and “personal slurring terms,” like ‘fatso’, and ‘jerk’, and ‘asshole’ (Hay 2013; Jeshion 2013b; Bolinger 2017; Bach 2018). 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

something general about DCs' semantics, or by something general about their pragmatic sociopolitical contexts of utterance, is something I likewise wish to leave open.

So I haven't told you what DCs are. But Williamson (to return to our paradigmatic example) doesn't define his target expressions either. Williamson offers the claim that slurs like 'Boche' add to the descriptive content of "neutral" expressions like 'German' by carrying an additional conventional implicature. This is theory, not definition. And to the extent that it's theory, it is supposed to be supported by the best available evidence about the use practices of competent speakers. Since most of us (again, presumably) do not use slurs, the best evidence we have as theorists about the use patterns of competent speakers comes from our observations of others. In Williamson's case, it comes apparently from Dummett's (1973) report of 'Boche'-users. In contrast, I believe that all of our ability to extend the foregoing list of DCs shows that all of us have the relevant kind of competence with DCs. We may not exhibit that competence by using paradigmatic slurs, but we exhibit it by using other terms on the list (or others I haven't mentioned). In this way, we can collect evidence about slurs—which, after all, are DCs, whatever else they are—in a more direct way than Williamson (and indeed nearly anyone in the philosophical literature) does.

Indeed, as ordinary competent speakers, we are liable to notice a trend among many of the DCs on my list. In particular, we are liable to notice that it's not just the worst, most paradigmatic slurs among them which seem to have what philosophers have called "neutral counterparts": most of the DCs on my list may be intuitively "paired up" with other, (typically) non-pejorative group or category expressions. Somewhat awkwardly, but again to avoid begging questions, I'll call these (putatively) more "neutral" group or category expressions non-pejorative associates, or NPAs:

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 Jeshion calls "all purpose put-downs", 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 (2013b: 236). Nevertheless, she maintains that paradigmatic slurs are distinctive in their relationships to so-called "neutral counterparts."

Derogatory Classifier	Non-Pejorative Associate	Derogatory Classifier	Non-Pejorative Associate
'n*gger'	'Black (person)' 'African American (person)'	'scab'	'strikebreaker' 'person who crosses picket lines'
'k*ke'	'Jewish (person)', 'Jew'	'gangbanger'	'gang member'
'ch*nk'	'Chinese (person)'	'pig'	'police officer' 'cop'
'oriental'	Asian (person)'	'jarhead'	'(U.S.) marine'
'curry muncher'	'South Asian (person)'	'shrink'	'psychiatrist'
'sp*c'	'Hispanic (person)'	'tech bro'	'(male) coder'
'wetback'	'Mexican (immigrant)'	'parasite'	'landlord'
'gringo'	'English speaker'	'hobo'	'homeless (person)' 'unhoused (person)'
'cracker'	'white (person)'	'junkie'	'illicit drug (ab)user' '(heroin, cocaine) drug addict'
'goy'	'non-Jewish (person)'	'pillhead'	'prescription drug (ab)user' 'opioid addict'
'Bible banger'	'(Evangelical) Christian'	'alchie'	'alcoholic'
'towelhead'	'Muslim (person)' 'Arab (person)'	'stoner'	'cannabis user', 'marijuana user' 'weed smoker'
'tr*nny'	'transgender' (person)	'frat bro'	'(male) fraternity member'
'f*ggot'	'gay (man)' 'homosexual (man)'	'gamer'	'person who plays video games'
'poof'	'gay (man)' 'homosexual (man)'	'jock'	'athlete' 'person who likes sports'
'd*ke'	'lesbian (woman)' 'homosexual (woman)'	'horse girl'	'girl who likes horses'
'rug muncher'	'lesbian (woman)' 'homosexual (woman)'	'ginger (kid)'	'redhead', 'person with red hair'
'breeder'	'straight (person)' 'heterosexual (person)'	'fatso'	'fat (person)'
'c*nt'	'woman' 'female (person)'	'trailer trash'	'person who lives in a trailer'
'r*tard'	'cognitively disabled (person)'	'townie'	'[town] native'
'cripple'/'crip'	'disabled (person)'	'tourist trap'	'tourist attraction'
'spaz'	'(disabled) person with a movement disorder'	'flyover state'	'U.S. state between coasts'
'boomer'	'Baby Boomer' 'person born 1946-1964'	'McMansion'	'(expensive) tract home'
'commie'	'communist'	'pill mill'	'pain clinic'
'tankie'	'(Stalinist) communist'	'rainbow corp'	'pro-LGBTQ company'

(Contd.)

Derogatory Classifier	Non-Pejorative Associate	Derogatory Classifier	Non-Pejorative Associate
'libtard'	'(American) liberal'	'breeder bar'	'straight bar' 'bar for non-LGBTQ people'
'woke' (n.)	(American) 'progressive'	'rag'	'newspaper'
'feminazi'	'feminist'	'gas guzzler'	'vehicle with low gas mileage'
'Bernie bro'	'male Bernie Sanders supporter'	'man cave'	'basement', 'den'
'Trumper'	'Donald Trump supporter'	'mom jeans'	'jeans a mom wears'
'anti-vaxxer'	'vaccine skeptic' 'person leery of vaccines'	'dad joke'	'joke told by a dad'
'treehugger'	'environmentalist'	'dad bod'	'body (type) of a dad'
		'capeshit movie'	'superhero movie' 'comic book movie'
		'chick flick'	'romantic comedy'

Call any such pair of intuitively-linked expressions a “DC/NPA pair”:

DC/NPA pair for any two group or category expressions x and y , if x and y are intuitively linked in meaning, but x is (typically) derogatory and y is (typically) not, then ' x '/' y ' is a DC/NPA pair.

Like examples of DCs, examples of intuitive DC/NPA pairs are easy to generate. Indeed, I suspect that, having now detected a pattern, the reader will find it easy to come up with more. This again is strong *prima facie* evidence that *DC/NPA pairs are a highly general linguistic phenomenon, admitting of a basically unifying linguistic explanation*. That there is a basic “DC/NPA relationship” which they all share in common is thus my default hypothesis:

generality all intuitive DC/NPA pairs in English are unified by a general linguistic relationship.

If GENERALITY 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.

But if GENERALITY is true, then this strongly suggests that our semantic theory of paradigmatic slurs should extend to DCs generally. And notably, almost no existing theory of slurs can be generalized in this way. For example, if Williamson's theory were to be generalized to all DCs, it would be the thesis that every DC has an NPA with which it shares its descriptive content, and from which it differs only in carrying an additional conventional implicature. Even already we should be skeptical of the this hypothesis, since it's hard to come up with NPAs for some DCs (like 'hillbilly', 'chav', and 'backwater'). But I'm not going to rest

my case on that. Let's suppose instead that Williamson's theory were to be generalized only to intuitive DC/NPA pairs. Abstracting to accommodate GENERALITY, this amounts to:

DC/NPA coextension for any DC and NPA intuitively linked in meaning, the linguistic relationship between them involves semantic coextension.

My argument against DC/NPA COEXTENSION is very simple: it's that not all chick flicks are romantic comedies, and not all romantic comedies are chick flicks. The rest of this paper consists entirely in generalizing and drawing morals.

The generalizing will proceed along two dimensions. The first, which I take up in §§3–8, is from the pair 'chick flick'/'romantic comedy' to other DC/NPA pairs. I begin by presenting a range of linguistic data which philosophers of language have, regrettably, hereto overlooked. Having thus begun to motivate the idea that DC/NPA pairs, including paradigmatic slur/"neutral counterpart" pairs, appear, *prima facie* at least, to pattern with 'chick flick'/'romantic comedy', I then propose an "overlap hypothesis" for 'chick flick'/'romantic comedy' in §4. In §5, I argue that extending this same kind of overlap hypothesis to other DC/NPA pairs, more than being merely suggested by the data presented in §4, also explains the apparently systematic behavior of DC/NPA pairs in certain peculiar kinds of conversational exchanges. Moreover, as I argue in §§6–8, it explains that behavior in a simpler and more unified way than salient alternative hypotheses.

The second dimension of generalization, which I will take up in §9, is from DC/NPA COEXTENSION and Williamson's conventional implicature view to nearly every other theory in the existing philosophical slurs literature. A newcomer to that literature may be surprised to discover that the theories on offer there, despite their manifest number and apparent diversity, almost universally presuppose that the thesis of this paper, and more broadly GENERALITY, is false. This is not because GENERALITY itself is rejected as a salient hypothesis, but because a foreclosing assumption about slurs and so-called "neutral counterpart" terms is treated as near-axiomatic. This assumption, which I will call the "JUST-ADD-BAD" assumption about so-called "neutral counterparts", is especially clear in Williamson's view; as we'll see in §9, a much wider set of theories of slurs make a similar mistake. In §10, I offer a diagnosis for why that is.

3. Data

My first charge against standard theories of slurs is methodological. For reasons I will speculate about later, philosophers have consistently overlooked huge swaths of the relevant empirical data. They have (unsurprisingly!) focused on

“paradigmatic” slurring expressions; but if GENERALITY is true, then we should expect the relevant pool of expressions to be much, much wider. In particular, we should expect to find patterns of use among not only competent users of paradigmatic slurs and their so-called “neutral counterparts”, but competent users of other DCs and NPAs as well. And the patterns we do find are striking indeed!

I will showcase two in particular. Both of these patterns involve explicitly *contrastive* uses of the relevant DCs and NPAs. In the first pattern, the relevant contrast is *evaluative*; users express differential attitudes of (comparable) approval or acceptance toward members of the relevant groups. In the second pattern, the relevant contrast is *predicative*; users express beliefs that one term applies but not the other.

3.1. *Contrastive Evaluations and Prescriptions*

Whatever the relationship is between DCs and NPAs, it is partially contrastive. This is to some extent true by stipulation, insofar as DCs are, to non-users at least, typically pejorative expressions, and NPAs (to non-users) typically are not. But beyond this perceived difference in (typical) pejorative force, it is also an empirical datum that DCs and NPAs are routinely *contrastively evaluated* by speakers who actually use them. The following kinds of evaluative judgments and prescriptions, for example, are made by competent users all the time:

- (8) From one Jew to another don't be a Kike.¹³
- (9) It's fine being black but don't be a nigger everywhere you go.¹⁴
- (10) 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.¹⁵
- (11) It's okay to be gay, it's not okay to be a faggot tho.¹⁶
- (12) Be a lesbian & not a dyke, theres a huge difference...¹⁷
- (13) omg i so did NOT call u a wetback i said mexican there is a big difference! trust me.¹⁸
- (14) It's okay to be liberal, but it's not okay to be a libtard.¹⁹
- (15) I'm OK with people who support Bernie. I do not like Bernie Bros.²⁰

13. Twitter, 11 Mar 2020. [[direct](#)] [[archive](#)]

14. Twitter, 23 May 2019. [[direct](#)] [[archive](#)]

15. Twitter, 9 Mar 2018. [[direct](#)] [[archive](#)]

16. Twitter, 13 Nov 2016. [[direct](#)] [[archive](#)]

17. Twitter, 28 Mar 2011. [[direct](#)] [[archive](#)]

18. Twitter, 11 Aug 2010 [[direct](#)] [[archive](#)]

19. Twitter, 2 Sep 2018. [[direct](#)] [[archive](#)]

20. Twitter, 20 May 2016. [[direct](#)] [[archive](#)]

- (16) You can be wary of Big Pharma and corrupt stuff going on [with vaccines] and unethical practices, but don't be an antivaxxer.²¹
- (17) I'm not saying fuck the police. I'm saying fuck the pigs. There's a difference between a police officer and a pig.²²
- (18) There's a difference between a Hobo, A junkie, and a Homeless person... Hobo=drunk, Junkie=thief, Homeless=deserve help.²³
- (19) There's a HUGE difference between living in a trailer and being trailer trash...²⁴
- (20) I'm an environmentalist, but I hate treehuggers.²⁵
- (21) I like redheads, but not gingers.²⁶
- (22) Don't mind people who smoke weed, HATE FUCKING STONERS!²⁷
- (23) There is a massive difference between people that play games and "gamers."²⁸
- (24) Act like a fraternity brother, and not a frat bro.²⁹
- (25) Be an athlete not a jock.³⁰
- (26) Even if you are a mom, it's not okay to wear mom jeans.³¹
- (27) Dads are hot but not dad bods.³²
- (28) Craft rooms, hobby rooms, office rooms, and dens are fine. But I draw the line at man caves.³³
- (29) There's a fundamental difference between superhero films and capeshit. Joker is a superhero film, MCU is capeshit.³⁴
- (30) I like romantic comedies. Not chick flicks. There's a huge difference.³⁵

This fact, along with the additional observation that among such examples there are several common patterns, is consistent with our working hypothesis about DC/NPA pairs. And supposing that that hypothesis is true, it places an empirical constraint on any adequate theory of the DC/NPA relationship: any such theory should plausibly explain not only why particular DCs and NPAs are intuitively

-
21. Twitter, 7 Dec 2020. [[direct](#)] [[archive](#)]
 22. Twitter, 29 Apr 2015. [[direct](#)] [[archive](#)]
 23. Twitter, 9 Jun 2012. [[direct](#)] [[archive](#)]
 24. Twitter, 26 Apr 2021. [[direct](#)] [[archive](#)]
 25. Twitter, 10 Jan 2019. [[direct](#)] [[archive](#)]
 26. Twitter, 29 Aug 2012. [[direct](#)] [[archive](#)]
 27. Twitter, 7 Nov 2013. [[direct](#)] [[archive](#)]
 28. Twitter, 17 Apr 2021. [[direct](#)] [[archive](#)]
 29. Twitter, 23 Mar 2013. [[direct](#)] [[archive](#)]
 30. Twitter, 8 Jan 2014. [[direct](#)] [[archive](#)]
 31. Twitter, 27 Dec 2012. [[direct](#)] [[archive](#)]
 32. Twitter, 13 Aug 2021. [[direct](#)] [[archive](#)]
 33. Twitter, 2 Sep 2010. [[direct](#)] [[archive](#)]
 34. Twitter, 14 Nov 2020. [[direct](#)] [[archive](#)]
 35. Twitter, 27 Jan 2012. [[direct](#)] [[archive](#)]

linked for competent speakers, but also why competent users of those expressions often explicitly contrast them in evaluative judgments and commands.

3.2. *Bidirectional Ascription Divergence*

In §2, I emphasized how competent users of ‘chick flick’ and ‘romantic comedy’ are prepared to grant that some sentences like (1) and (2) are (or in principle could be) true:

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

In such cases, competent users take ‘chick flick’ and ‘romantic comedy’ to come apart in both directions. But this isn’t just true of ‘chick flick’/‘romantic comedy’; competent users of many, many other DCNPA pairs use them in a structurally analogous way.

Examples of such “bidirectional divergence” in DC/NPA ascriptions are exceedingly easy to come by in “real life” conversational contexts. My goal in this section is merely to showcase their apparent systematicity and range:

‘antivaxxer’/‘person cautious of vaccines’

- (31) You don’t have to be against all vaccines to be an antivaxxer. if you’re cautious of a vaccine that’s proven safe and effective but not of the infectious disease which is more likely to kill and hospitalise, then you’re an antivaxxer.³⁶
- (32) As someone with an autoimmune condition I’m leery of vaccines (I’m not an ‘antivaxxer’ just well aware my body can go haywire after getting a jab).³⁷

‘Bernie bro’/‘male Bernie Sanders supporter’

- (33) Just because you’re a Bernie supporter doesn’t mean you’re a Bernie bro.³⁸
- (34) You don’t have to be a man to be a Bernie Bro, you just have to act like a hooligan.³⁹

36. Twitter, 24 Sept 2021. [[direct](#)] [[archive](#)]

37. Twitter, 13 Mar 2021 [[direct](#)] [[archive](#)]

38. Twitter, 23 Feb 2020. [[direct](#)] [[archive](#)]

39. Twitter, 17 Feb 2020. [[direct](#)] [[archive](#)]

'boomer'/'Baby Boomer'[(American) person born 1946–1964]

- (35) I differentiate it this way: Baby-boomer is the generation. Boomer (by itself) is an insult to connote a state of mind. You're a baby-boomer, but not a boomer.⁴⁰
- (36) Mayo Pete is a boomer at 37. Does this help everyone understand that boomer isn't just an age thing?⁴¹

'breeder'/'straight (person)'

- (37) I will never understand breeders. No not two straight people having sex, I mean people who actively choose to reproduce. Literally why would you ever? You're giving up your whole life away for a measly little spawn.⁴²
- (38) Any queer who has biological children is still a breeder.⁴³

'capeshit (movie)'/ 'superhero movie'

- (39) The Dark Knight is the rare instance of a superhero movie that isn't capeshit.⁴⁴
- (40) It's a sad fact to come to terms with but Star Wars is "capeshit".⁴⁵

'ch*nk'/'Chinese (person)' [note: first example also contains the N-word]

- (41) Not all Chinese people are chinks. Just like not all black people are niggers.⁴⁶
- (42) I hate those stupid Vietnamese people. Stupid chinks always torturing animals.⁴⁷

'c*nt'/'woman'

- (43) Not all women are cunts. Term is reserved for only the deserving.⁴⁸
- (44) Men can be cunts too. Just as they can be pussies.⁴⁹

40. Twitter, 6 Apr 2020. [[direct](#)] [[archive](#)]

41. Twitter, 19 Nov 2019. [[direct](#)] [[archive](#)]

42. Twitter, 9 Apr 2020. [[direct](#)] [[archive](#)]

43. Twitter, 3 May 2020. [[direct](#)] [[archive](#)]

44. Twitter, 21 Apr 2021. [[direct](#)] [[archive](#)]

45. Twitter, 20 Feb 2020. [[direct](#)] [[archive](#)]

46. Forum comment, 29 July 2017. [[direct](#)] [[archive](#)]

47. Twitter, 5 June 2015. [[direct](#)] [[archive](#)]

48. Twitter, 11 Dec 2016. [[direct](#)] [[archive](#)]

49. Twitter, 11 Dec 2016. [[direct](#)] [[archive](#)]

'dad joke'/'joke told by a dad'

- (45) Being a dad with “actually funny jokes and not just dad jokes” might be the best feedback you can get from your kid.⁵⁰
- (46) I don't understand why Cory Booker tells so many Dad jokes when he's not a Dad.⁵¹

'dad bod'/'body [type] had by a dad'

- (47) You don't have to be a dad to have a dad bod.⁵²
- (48) Lesley and I saw probably the hottest dad ever without a dad bod...god bless.⁵³

'd*ke'/'lesbian (person)'

- (49) I always wanted a lesbian friend, lesbian not a dyke.⁵⁴
- (50) Not all dykes are lesbians. I got a cousin who's a dyke but she has a husband.⁵⁵

'f*ggot'/'gay (person)'

- (51) Just remember honey, there are good gays, and bad gays. There are gays and then there are faggots.⁵⁶
- (52) Whoever decided to allow transgender faggots in the military is a retard.⁵⁷

'fatso'/'fat person'

- (53) You're a fatso. You're skinny but you eat a lot.⁵⁸
- (54) “You're fat but not so fat to be called a fatso” My little brother is so kind to me.⁵⁹

'gamer'/'person who plays video games'

- (55) I deliberately avoid identifying as a gamer because of the toxic associa-

50. Twitter, 2 Apr 2019. [[direct](#)] [[archive](#)]

51. Twitter, 15 Oct 2019. [[direct](#)] [[archive](#)]

52. Twitter, 1 May 2021. [[direct](#)] [[archive](#)]

53. Twitter, 2 Jul 2018. [[direct](#)] [[archive](#)]

54. Twitter, 28 June 2013. [[direct](#)] [[archive](#)]

55. Twitter, 9 Nov 2019. [[direct](#)] [[archive](#)]

56. Twitter, 17 Nov 2017. [[direct](#)] [[archive](#)]

57. Twitter, 30 Jun 2016. [[direct](#)] [[archive](#)]

58. Twitter, 27 Apr 2017 [[direct](#)] [[archive](#)]

59. Twitter, 27 Dec 2012. [[direct](#)] [[archive](#)]

- tions of gamer culture. I am a game player, for sure, but not a gamer.⁶⁰
- (56) Typical gamer. Doesn't even play the game, just wanks off at the pretty girls.⁶¹

'ginger'/'redhead'

- (57) You stan a ginger I don't wanna hear it. And before you say Shanks is one too no he's just a redhead but not a ginger.⁶²
- (58) She's a ginger, not the redhead ginger though. But still a ginger. And they say gingers have no soul.⁶³

'horse girl'/'girl who has horses'

- (59) Alright she has horses but she's not a 'horse girl'.⁶⁴
- (60) You don't have to own or ride a horse to be a horse girl. It's a certain je ne sais quoi you have.⁶⁵

'junkie'/'(heroin) drug addict'

- (61) Getting addicted does in fact make you an addict! But it doesn't make you a junkie by any means.⁶⁶
- (62) If you snort cocaine from time to time you're still a junkie I don't care.⁶⁷

'k*ke'/'Jew(ish)'

- (63) Maybe, but he isn't a kike or have a socialist agenda like 99% of the Jews.⁶⁸
- (64) The Church at Old Jewish Schildesche wanted me dead you kike Catholics.⁶⁹

'libtard'/'liberal'

- (65) 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.⁷⁰

60. Twitter, 23 Jun 2019. [[direct](#)] [[archive](#)]

61. Twitter, 30 May 2014. [[direct](#)] [[archive](#)]

62. Twitter, 7 Mar 2022. [[direct](#)] [[archive](#)]

63. Twitter, 13 Sep 2012. [[direct](#)] [[archive](#)]

64. Twitter, 30 Jan 2020. [[direct](#)] [[archive](#)]

65. Twitter, 11 Apr 2020. [[direct](#)] [[archive](#)]

66. Twitter, 19 Mar 2019. [[direct](#)] [[archive](#)]

67. Twitter, 12 Oct 2018. [[direct](#)] [[archive](#)]

68. Twitter, 20 Nov 2016 [[direct](#)] [[archive](#)]

69. Twitter, 16 Feb 2018 [[direct](#)] [[archive](#)]

70. Twitter, 20 Aug 2018. [[direct](#)] [[archive](#)]

(66) *About someone wearing a communist shirt:* Maybe not the exact definition of liberal but still a libtard.⁷¹

'mom jeans'/'jeans worn by a mom'

(67) My mom always wearing jeans. Thank god she doesn't wear mom jeans!⁷²

(68) Lucky Obama doesn't have a son! How would you like to grow up with a dad who throws like a girl & wears mom jeans?⁷³

'pig'/'police officer'['cop']

(69) Not everyone who consumes marijuana is a violent criminal/bum, just like not every police officer is a trigger happy pig.⁷⁴

(70) Judges, prosecutors, guards, and military [aren't police officers but] are still pigs.⁷⁵

'rag'/'newspaper'

(71) The Guardian is a lefty liberal newspaper but not a rag. I still regard it one of British media's best in news reporting and analysis.⁷⁶

(72) I've had enough of this rag magazine [*Rolling Stone*]. It used to be respectable but now it's no better than a supermarket rag paper at the checkout.⁷⁷

'stoner'/'person who smokes weed'

(73) I want a roommate that smokes weed but isn't a stoner.⁷⁸

(74) You can be a stoner without smoking weed.⁷⁹

'trailer trash'/'person who lives in a trailer'

(75) "Nomadland" showed us that just because you live in a trailer, you're not necessarily trailer trash. Similarly, Marjorie Taylor Greene lives in a house.⁸⁰

71. Twitter, 11 July 2022. [[direct](#)] [[archive](#)]

72. Twitter, 30 Mar 2012 [[direct](#)] [[archive](#)]

73. Twitter, 15 Jul 2014. [[direct](#)] [[archive](#)]

74. Twitter, 29 Nov 2019. [[direct](#)] [[archive](#)]

75. Twitter, 6 Jun 2020. [[direct](#)] [[archive](#)]

76. Twitter, 13 Nov 2018. [[direct](#)] [[archive](#)]

77. Twitter, 9 Nov 2022. [[direct](#)] [[archive](#)]

78. Twitter, 31 Jan 2020. [[direct](#)] [[archive](#)]

79. Twitter, 7 Apr 2020. [[direct](#)] [[archive](#)]

80. Twitter, 25 May 2021 [[direct](#)] [[archive](#)]

The structural pattern, here, is striking—as far as it goes. But while such a pattern is consistent with my proposal—viz., that intuitive DC/NPA exhibit the same basic linguistic relationship as ‘chick flick’/‘romantic comedy’—it certainly does not entail it. Much more needs to be said to motivate this idea, especially for DC/NPA pairs involving paradigmatic slurs. This will be my goal in the next two sections.

4. ‘Chick Flick’/‘Romantic Comedy’: The Case for Overlap

4.1. An Obvious First Thought


The ubiquity, range, and apparent systematicity of the examples just surveyed lends *prima facie* support to the default hypothesis proposed in §2:

generality all intuitive DC/NPA pairs in English are unified by a general linguistic relationship.


More specifically, we seem to have evidence that DC/NPA pairs pattern with, and hence are basically assimilable to, ‘chick flick’/‘romantic comedy’. So whatever the relationship is between ‘chick flick’ and ‘romantic comedy’, we have *prima facie* reason to assume that an analogous relationship obtains for all other DC/NPA pairs.

What, then, is the nature of the relationship between ‘chick flick’ and ‘romantic comedy’?

If we weren’t ourselves competent speakers, and were to take our cue from the philosophical slurs literature, we might suppose that, extensionally, there are two possibilities for ‘chick flick’ and ‘romantic comedy’: either ‘chick flick’ refers to a *proper subset* of the things which ‘romantic comedy’ refers to; or ‘chick flick’ and ‘romantic comedy’ pick out *exactly the same extension*:




restricted extension. All chick flicks are romantic comedies, but not all romantic comedies are chick flicks.



coextension. All chick flicks are romantic comedies, and all romantic comedies are chick flicks.

As competent speakers, though, we know that both of these hypotheses are obviously nonstarters. Indeed, the obvious first thought is that, extensionally, ‘chick flick’ and ‘romantic comedy’ *generally*—but not completely—overlap:



overlapping extensions. In general, most chick flicks are romantic comedies, and most romantic comedies are chick flicks.

Moreover, the “obviousness” of this “obvious first thought” seems obviously not an accident. While competent users will grant that “chick flicks” and “romantic comedies” can, sometimes, fall in the margins of the Venn diagram, they expect them to fall in the middle. For competent speakers, ‘chick flick’ and ‘romantic comedy’ are prototypically directed: they are conceptually associated with most if not all of the same stereotypes, beliefs, and evaluative attitudes.

This, too, is a kind of “overlap”: call it *stereotype overlap*. Together, the notions of stereotype overlap and extensional overlap (hereafter “S-overlap” and “E-overlap”, respectively) give us an “obvious first answer” to our question:

overlap hypothesis the intuitive linguistic relationship between ‘chick flick’ and ‘romantic comedy’ substantial degrees of overlap in

- (a) semantic extension (E-overlap), and
- (b) the stereotypes, beliefs, and evaluative attitudes associated with each expression in the minds of competent speakers (S-overlap).

OVERLAP HYPOTHESIS has immediate intuitive appeal. But we need not rely solely on intuition to think it’s on the right track. The ultimate test of plausibility for any of the candidate pictures of the target relationship (restricted extension, coextension, overlap, etc.) is that they have to make sense of both the phenomenon I have been emphasizing so far (explicitly contrastive uses) and cases where competent users seem, at least, to treat the target expressions as synonyms. It would be ideal, then, for us to find a test case in which speakers seem to do both of these things. And indeed, there are such cases—viz., exchanges like (76):

- (76) A: The new Ghostbusters is a chick flick.
- B: But it’s an action movie, not a romantic comedy!

A: Whatever, it's still a chick flick.

In the next section, I will argue that OVERLAP HYPOTHESIS with some plausible assumptions about common ground, best explains this exceedingly familiar but, through the lens of the philosophical slurs literature, somewhat peculiar kind of exchange. Having made this case, I will then show that the same basic story, just as GENERALITY would have us predict, successfully generalizes to a range of other DC/NPA pairs.

4.2. A Default Presumption against Exceptional Cases

When competent users of 'chick flick' and 'romantic comedy' witness exchanges like (76), they effortlessly understand them. But what explains this, exactly? Why should B's utterance in (76) 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, in responding the way he does, has taken something for granted about the relationship between chick flicks and romantic comedies. In particular, B seems to have presupposed that 'chick flick' applies to something only if 'romantic comedy' does. Thus, as he believes that *Ghostbusters* isn't a romantic comedy, he tries to correct A's application of 'chick flick' to the film.


On reflection, B would likely reject this presupposition about chick flicks and romantic comedies. If asked to provide counterexamples like (1) and (2), he probably could:

- (1) The new *Ghostbusters* is chick flick, but it isn't a romantic comedy.
- (2) *Silver Linings Playbook* is a romantic comedy, but it isn't a chick flick.

Still, he would probably have to think about it a little. For when B, like most competent users, ordinarily hears⁸¹ the words 'chick flick' and 'romantic comedy', movies like *Ghostbusters* and *Silver Linings Playbook* are probably not what first comes to mind. What likely does first come to mind 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 competent users are concerned). For competent users, the most prototypical examples of chick flicks are

81. For 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.


also prototypical examples of romantic comedies. Less prototypical examples are, well, atypical: if cases like (1) and (2) obtain, they do so only at the margins:




overlapping extensions. In general, most chick flicks are romantic comedies, and most romantic comedies are chick flicks.

Indeed, competent users’ semantic commitment to substantial *E*-overlap between ‘chick flick’ and ‘romantic comedy’ is plausibly underwritten by the substantial *S*-overlap in their representational beliefs. And plausibly, this in turn has certain socio-pragmatic upshots.

In the first place, we should expect a semantic commitment to substantial *E*-overlap (underwritten by substantial *S*-overlap) to inform competent users’ default expectations about which sorts of things they will encounter in the world. But we should also, crucially, expect it to inform their default conversational expectations of one another. If competent users believe that the (vast) majority of “romantic comedies” they encounter in the world will be “chick flicks”, and that the (vast) majority of “chick flicks” they encounter will be “romantic comedies”, then it is plausible that they will believe of conversations with one another that possible exceptions will be “bracketed” by default. It is plausible, in other words, that in normal conversational contexts, competent users expect one another to proceed as if the *E*-overlap between ‘chick flick’ and ‘romantic comedy’ is complete:



-->



general (incomplete) *E*-overlap. Most chick flicks are romantic comedies, and most romantic comedies are chick flicks.

complete *E*-overlap (presumption). All chick flicks are romantic comedies, and all romantic comedies are chick flicks.

Such a default presumption against exceptional cases, underwritten by a shared commitment to substantial *E*-overlap, can explain some important data about exchanges like (76):

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

First, it can explain why B replies the way he does. Second, it can explain why, despite B's replying the way that he does, A's rejoinder is felicitous, and why the exchange in general is so predictable and familiar.

Initially, at least, B is confused by A's assertion. But this is exactly what we should expect, if B expects a presumption against exceptional cases to operative by default, and A to say so if it's not. Indeed, A could have cancelled this presumption for B, had he wanted to or thought to. In particular, he could have said (1):

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

But A doesn't say (1); he says (77):

- (77) The new *Ghostbusters* is a chick flick.

And this, together with the presumption that all and only chick flicks are romantic comedies, immediately entails (78)—which B rejects:

- (78) The new *Ghostbusters* is a romantic comedy.

Hence his attempted correction: "But it's an action movie, not a romantic comedy!"

Of course, this is not the only possible explanation for B's attempt to correct A. We might hypothesize, for example, that what B is attempting to correct in (76) is some obvious misuseage which he takes A to have committed—one which would throw into question A's basic competence with 'chick flick'. But we have little reason to assume this. A need not have made a semantic error for his conversational move to be suboptimal; B's complaint could just as easily be pragmatic. By choosing to say (77) when he could have said (1), and neglecting (*ex hypothesi*) to cancel the presumption that all and only chick flicks are romantic comedies, A has thus predictably and avoidably introduced confusion into their exchange.

Moreover, if A's use of 'chick flick' in (76) was indeed an obvious semantic mistake—the kind that would make us question his basic competence—then it would be hard to explain why his rejoinder is (a) so immediately comprehensible, and (b) not conversation-stopping.

There is no intuitive sense that A, in rejoining the way he does, is being somehow uncooperative, or failing to give B uptake. On the contrary, A intuitively

concedes B's rejection of 'romantic comedy' as applied to the new *Ghostbusters*—this is the intuitive content of his 'whatever'. Indeed, the casual dismissiveness of that same 'whatever' suggests that A is unsurprised by B's confusion. This makes sense, if A is a competent user who normally shares a presumption against exceptional cases. That he thinks this is an exceptional case explains why, having thus conceded the rejection of the NPA, he proceeds to double down on the DC. And that such alleged marginal cases are, from the perspective of competent users, possible, explains why his doing so is felicitous. B may not agree with A that the new *Ghostbusters* is a "chick flick," properly so-called; the two of them may go on to argue about it. But it would be strange, and indeed inappropriate, for B to simply throw up his hands at A's rejoinder in (79) and conclude he's speaking to an incompetent troll.⁸²

Now, recall DC/NPA COEXTENSION, which we abstracted from Williamson's (2009) conventional implicature view of paradigmatic slurs and generalized for all DC/NPA pairs:

DC/NPA coextension for any DC and NPA intuitively linked in meaning, the linguistic relationship between them involves semantic coextension.

DC/NPA COEXTENSION predicts that 'chick flick' and 'romantic comedy', if they indeed constitute an intuitive DC/NPA pair, are truth-conditionally equivalent.

I submit that this prediction is false on its face. But suppose that we thought it was at least *prima facie* plausible enough to test. An initial point in its favor is that it in turn predicts something importantly true about exchanges like (76)—viz., that competent users will conversationally presuppose that all and only chick flicks are romantic comedies. This is in fact what B seems to do!

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

But DC/NPA COEXTENSION also predicts much more than this—indeed, much too much. For if DC/NPA COEXTENSION were true, then we should also expect that the expressions 'chick flick' and 'romantic comedy' will, from the perspective of

82. Some readers, their hackles raised at this point by the repetition involved in (76), may be inclined to wonder if the felicity of A's rejoinder is due to this just being another "Frege Puzzle" case. But this is implausible. For (76) to involve "Frege Puzzle"-style confusion, it would require not only that 'chick flick' and 'romantic comedy' be co-extensional (which I submit is obviously false), but that speakers A and B don't know this fact. And any theory that posits such basic incompetence on the part of competent speakers is not worth considering.

people who actually use them, be intersubstitutable in their normal, literal uses. We should expect, in other words, (76) and (79) to be equally comprehensible:

(79) A: The new *Ghostbusters* is a romantic comedy.

B: But it's an action movie, not a romantic comedy!

A: #Whatever, it's still a romantic comedy.

(79), however, is not comprehensible at all! Whereas, before, A *feliculously* doubled-down in the face of B's attempted correction, his rejoinder here is bizarrely uncooperative. And this actually should not be surprising, given what we've supposed already about A's rejoinder in (76)—viz., that the intuitive conversational function of his qualificational 'whatever' is to signal concession of B's point that the new *Ghostbusters* isn't a romantic comedy. But this is just the negation of A's original claim. A cannot felicitously concede that the new *Ghostbusters* isn't a romantic comedy, while also stubbornly insisting that it is.

This pronounced asymmetry in comprehensibility and felicitousness between (76) and (79) is a datum about 'chick flick' and 'romantic comedy' which we might not have discovered, had we not thought to try fitting those expressions into Williamson's mold. But it is a datum which any good theory of that pair should fit, all the same. And happily, OVERLAP HYPOTHESIS can already do so, without any fancy theoretical footwork or added machinery:

overlap hypothesis the intuitive linguistic relationship between 'chick flick' and 'romantic comedy' involves substantial degrees of overlap in

- (a) semantic extension (*E*-overlap), and
- (b) the stereotypes, beliefs, and evaluative attitudes associated with each expression in the minds of competent speakers (*S*-overlap).

OVERLAP HYPOTHESIS commits us to *substantial E*-overlap between 'chick flick' and 'romantic comedy'; but, crucially, it does not commit us to *complete E*-overlap. Indeed, this was the point: to capture the intuitive idea (from the perspective of competent users) that movies can be chick flicks without being romantic comedies, and can be romantic comedies without being chick flicks (though we may debate about which ones). But no movie can, in any literal sense, be a romantic comedy without being a romantic comedy. That a speaker cannot "get away" with denying this latter claim, but may give voice to the preceding one with his reputation as competent speaker intact, is not, on the present theory, any surprise at all.

In sum, then, the "obvious first thought" to have about 'chick flick' and 'romantic comedy' is that their relationship involves substantial *E*- and *S*-over-

lap; and this OVERLAP HYPOTHESIS, together with some very basic Stalnakerian assumptions about common ground, can successfully account for and predict several important data by which any adequate theory of that pair must be constrained:

- (a) that competent users of ‘chick flick’ and ‘romantic comedy’ take those expressions to be tightly intuitively related, both in extension and in the stereotypes and attitudes associated with them;
- (b) that, despite this fact, competent users sometimes apply ‘chick flick’ and ‘romantic comedy’ to different things;
- (c) that exchanges like (76) are, to competent users of ‘chick flick’ and ‘romantic comedy’, familiar and immediately comprehensible;
- (d) that B’s reply in (76) seems to presuppose that all and only chick flicks are romantic comedies (viz., that *E*-overlap between ‘chick flick’ and ‘romantic comedy’ is complete);
- (e) that despite this fact, A’s rejoinder in (76) is not only felicitous, but is consistent with, and indeed actually suggestive of, basic competence with the relevant expressions; and
- (f) that this marks a clear asymmetry with (79), where the exchange in general is defective, and A’s rejoinder in particular is bizarre and uncooperative.

Moreover, the present proposal explains why the foregoing would all be true, by explaining why, in normal conversational contexts, competent users would share a default presumption against exceptional cases.

This is strong evidence in favor of OVERLAP HYPOTHESIS. It is also, if the data in §3 is any indication that GENERALITY is true, evidence for a generalize version for *all* DC/NPA pairs:

DC/NPA overlap for any DC and NPA intuitively linked in meaning, the linguistic relationship between them involves substantial degrees of overlap in

- (a) semantic extension (*E*-overlap), and
- (b) the stereotypes, beliefs, and evaluative attitudes associated with each expression in the minds of competent speakers (*S*-overlap)

As we’ll see, DC/NPA OVERLAP is not quite the thesis I will ultimately land on. But, like its ‘chick flick’/‘romantic comedy’ analogue, it is the “obvious first thought” to have, based on what we have seen so far; sticking with it for now will make it easier to fix ideas. It also, I think, is not very far from the truth.

5. Generalizing to Other DC/NPA Pairs

In the previous section, I began by suggesting that whatever theory we accept about the relationship between ‘chick flick’ and ‘romantic comedy’, we have good reason to entertain about DC/NPA pairs in general. In this section, I will provide some additional arguments for this claim. In particular, I will argue that (a) DC/NPA pairs pattern systematically with ‘chick flick’/‘romantic comedy’ in exchanges like (76); and (b) such patterns are plausibly explained by the same pragmatic mechanism which I proposed for (76), whereby competent users proceed conversationally by default as if *E*-overlap is complete.

That story, recall, involved a default conversational presumption against exceptional cases among competent users, underwritten by a shared expectation that most cases will be typical cases; viz., cases wherein ‘chick flick’ and ‘romantic comedy’ do not come apart. This explains why, in (76), B attempts to correct A’s use of ‘chick flick’ by rejecting the aptness of ‘romantic comedy’, but A nevertheless felicitously doubles down:

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

This marked a clear contrast with (79), where B instead attempts to correct A’s use of ‘romantic comedy’, and A infelicitously doubles down:

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

Call exchanges like (76) *DC-corrections* and exchanges like (79) *NPA-corrections*.

In this section, I will show that the same asymmetry between DC- and NPA-corrections obtains for other DC/NPA pairs—including more orthodox pairs of paradigmatic slurs and “neutral counterpart” terms—and that in each case this asymmetry is plausibly explained by the same kind of overlap hypothesis proposed for ‘chick flick’/‘romantic comedy’ in §4. Moreover, I will argue that this is a better explanation than DC/NPA COEXTENSION, our toy Williamsonian foil.

Philosophers who are drawn to the Williamsonian view are likely to think that slurs and their neutral counterparts are importantly disanalogous from DC/NPA pairs like ‘chick flick’/‘romantic comedy’. So part of my goal will be to argue that they are not so different after all. Proponents of the Williamsonian view might point to any of three purported disanalogies: (i) difference in presence of derogatory force or intent, (ii) difference in degree of hypothesized overlap, and

(iii) difference in target kind. But as I will show, these are not *general* differences between slur/“neutral counterpart” pairs and other DC/NPA pairs. For each of (i), (ii), and (iii), there are some non-orthodox DC/NPA pairs which are not different from slur/“neutral counterpart” pairs in the relevant respect. Thus, we have every reason to pair that the basic DC/NPA relationship is *perfectly general*, making no special exception for more orthodox pairs of slurs “neutral counterparts”. A semantic theory of slurs ought to be generalizable to all DC/NPA pairs. DC/NPA COEXTENSION is generalized in this way, but the Williamsonian view is not and cannot be generalized in this way.

5.1. Charge: Difference in Presence of Derogatory Force or Intent

Suppose we grant OVERLAP HYPOTHESIS for ‘chick flick’/‘romantic comedy’. I claim that, on the basis of parity, we have *pro tanto* reason to assume other DC/NPA pairs work the same way. But it might be worried that this *pro tanto* reason indeed isn’t much of a reason at all, as ‘chick flick’ “isn’t really” a DC. It has been occasionally put to me, for example, that the genre term—qua genre term—isn’t obviously derogatory; or if it is now, that it wasn’t intended to be that way.

Setting aside that many DCs (including some of the most paradigmatic! Consider the onceclinical ‘R-word’) were in fact not originally meant as derogatives, let us grant that either or both of the above hypotheses about ‘chick flick’ is true. DC/NPA OVERLAP is equally as plausible for the genre term ‘capeshit (movie)’, an explicitly derogatory expression, coined to be derogatory. Initially popularized on the internet forum 4chan, ‘capeshit (movie)’ (sometimes shortened to ‘cape movie’ or ‘cape flick’) is a derogatory category expression with an intuitive NPA (‘superhero movie’, or ‘comic book movie’). Its online usage over the last decade has become increasingly widespread, as big-budget superhero movies like the *Avengers* franchise have proliferated.⁸³ And indeed, the connection between it, as a DC, and the expression ‘superhero movie’ (or ‘comic book movie’) is, for competent users, extremely tight. Still, competent users will entertain, if not in every particular instance accept, claims to the effect something is a “superhero movie” but not “capeshit (movie)”, or vice versa:

- (39) The Dark Knight is the rare instance of a superhero movie that isn’t capeshit.
- (40) It’s a sad fact to come to terms with but Star Wars [though not a superhero movie] is “capeshit”.

83. According to the website [KnowYourMeme](http://KnowYourMeme.com), the expression first appeared on 4chan in 2009.

Like ‘chick flick’/‘romantic comedy’, this is a case where an overlap hypothesis has immediate plausibility. I would even submit that it is obvious, in this case, that familiarity with the relevant stereotypes—and the negative attitudes typically attached to those stereotypes among relevant speakers—is essential for basic competence with sentences like (39) and (40), and likewise for DC-corrections like (80) and (81):

(80) A: The Star Wars movies are capeshit movies.

B: But they’re not superhero movies!

A: Whatever, they’re still capeshit movies.

(81) A: Joker isn’t a capeshit movie.

B. But it’s a comic book movie!

A: Whatever, it’s still not capeshit movie.

These DC-corrections, just like (76) are immediately comprehensible to speakers with basic competence with the relevant expressions. But to have the requisite competence, it is not enough to simply know that DC ‘capeshit (movie)’ is associated with ‘superhero movie’ and ‘comic book movie’. *It may seem sufficient in most contexts where ‘capeshit (movie)’ is used, because in most such contexts, the films at issue are also ones to which ‘superhero movie’ and ‘comic book movie’ obviously apply.* But in order to understand exchanges like (80) and (81), one must have a clearer sense of what users of the term ‘capeshit (movie)’ intend to be targeting about the kinds of movies they are expressing contempt for.

Often, this information is left tacit; but occasionally, speakers will make their meaning more explicit:

(82) A: The Star Wars movies are capeshit movies.

B: But they’re sci fi fantasy movies, not superhero movies!

A: Whatever, they’re still capeshit movies. They use their content to sell toys and merchandise to man-babies, and actual babies. Therefore since we’re all adults they should not be talked about in a public space.⁸⁴

(83) A: Joker isn’t a capeshit movie.

B: But it’s a comic book movie!

A: Whatever, it’s still not a capeshit movie. Capeshit movies are usually the same rehash “save the planet, beat the baddie.” Joker was babies

84. Twitter, 23 Dec 2019. [[archive](#)]

first Taxi Driver. But at least it was different than most MCU [Marvel Cinematic Universe] & DC films.⁸⁵

The key stereotypes, here, are features like being about superheroes/characters who have superpowers and wear capes, being about characters based on comic books, being cinematically generic/unoriginal, being cinematically simple-minded/brain-rotting, having lots of prequels and sequels, being over-reliant on CGI effects, being childish, and catering to (the ignorant) masses.⁸⁶

Crucially, these are stereotypes which competent users associate not just with the DC ‘capeshit (movie)’, but also with the NPAs ‘superhero movie’ and ‘comic book movie’. In this sense, these latter expressions (like ‘romantic comedy’ before and, I submit, NPAs more generally) are not really “neutral” at all, from the perspective of speakers who use the pejorative DC. An overlap thesis grounded in S-overlap predicts this, as it posits overlap not just in user’s associated stereotypes and descriptive beliefs, but also in their negative evaluative attitudes. Still, when they want to, competent users can—and do—recover more a “technical” sense of ‘superhero movie’ (or ‘comic book movie’), where these just mean “movie about superheroes” (or “movie based on comic books”). Thus the speaker of (39) defends *The Dark Knight*:

(39) The Dark Knight is the rare instance of a superhero movie that isn’t capeshit.

Notably, this is also what intuitively happens with ‘chick flick’/‘romantic comedy’ in cases like (2), where a ‘chick flick’-user “defends” a movie which they acknowledge is technically a “romantic comedy”:

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

And indeed, just as with ‘chick flick’/‘romantic comedy’, there is an clear asymmetry in comprehensibility and felicitousness between the above DC-corrections and (84):

(84) A: The Star Wars movies are superhero movies.
 B: But they’re sci fi fantasy movies, not superhero movies!
 A: #Whatever, they’re still superhero movies.

85. Twitter, 26 Oct 2019. [direct] [archive]

86. In the course of researching for this paper, I discovered that posters on Twitter and especially 4chan often refer to franchises like Marvel’s *The Avengers* as “a cancer”.

If we find that we can “make sense” of this exchange at all, it is because we force it to, by reading A as meaning something metaphorical by ‘superhero movie’, such that his assertion is consistent with the *Star Wars* movies not (actually) being superhero movies. But it is the basic relationship between DCs and NPAs, in their literal uses, which we are trying to theorize, here. And if we instead force a literal reading of A’s claims in (84), then his rejoinder obviously crashes:⁸⁷

(85) A: The *Star Wars* movies are superhero movies.

B: But they’re sci fi fantasy movies, not movies about superheroes!

A: #Whatever, they’re still superhero movies. They’re movies about superheroes.

And as with ‘chick flick’/‘romantic comedy’, this asymmetry is plausibly explained by competent users’ mutual commitment to substantial extensional overlap (grounded in substantial stereotype overlap), but not complete extensional overlap between the DC and the NPA.

5.2. Charge: *Difference in Degree of Overlap*

Just as nothing turns on the relevant DCs being “mild” or plausibly nonderogative in origin, nothing in the present “overlap” model requires that the *degree* of overlap be the same for all DC/NPA pairs. It requires only that there be *enough* E- and S-overlap to allow for a presumption of complete overlap among ordinary speakers. DC/NPA OVERLAP is thus more permissive than it might initially have seemed: it is compatible with a range of *degrees* of the relevant overlap, including ones where the relevant DC/NPA pairs are, intuitively, less extensionally and conceptually tied than ‘chick flick’ and ‘romantic comedy’.

To illustrate, I will now present two cases where the hypothesized overlap, and speakers’ expectations of overlap, profiles differently, viz., ‘mom jeans’/‘jeans worn by a mom’ and ‘dad joke’/‘joke told by a dad’. Like ‘chick flick’ and ‘capeshit (movie)’, ‘mom jeans’ and ‘dad joke’ are intuitive (if extremely mild) DCs: they are nominalized group or category expressions with established pejorative uses in certain speech communities. They also have intuitive NPAs: ‘jeans belonging to a mom’ and ‘joke told by a dad’, respectively—where the former picks out all and only the jeans worn by people who are moms, and the latter picks out all and only the jokes told by people who are dads. And these

87. I will return to the issue of possible metaphorical readings in §6.

intuitive relationships, like the one between ‘chick flick’ and ‘romantic comedy’ and ‘capeshit (movie) and ‘superhero movie’, seem to involve overlap.

Indeed, as with those previous cases, I submit that it is obvious that these other relationships involve overlap. Competent users of ‘mom jeans’ are obviously committed to substantial but incomplete *E*-overlap: while they think that some of the jeans worn by moms are “mom jeans” (viz., the worst kind), they also believe that one need not actually be a mom to commit this particular fashion faux pas:

- (67) My mom always wearing jeans. Thank god she doesn’t wear mom jeans!⁸⁸
 (68) Lucky Obama doesn’t have a son! How would you like to grow up with a dad who throws like a girl & wears mom jeans?⁸⁹

Likewise, competent users of ‘dad joke’ obviously believe that while some of the jokes told by people who are dads are “dad jokes”, being a dad is neither necessary nor sufficient to be “guilty” of the jokes in question:

- (45) Being a dad with “actually funny jokes and not just dad jokes” might be the best feedback you can get from your kid.⁹⁰
 (46) I don’t understand why Cory Booker tells so many Dad jokes when he’s not a Dad.⁹¹

And this commitment among competent users seems obviously grounded in *S*-overlap between associated stereotypes and attitudes.

In the case of ‘mom jeans’/‘jeans belonging to a mom’, the relevant stereotypes include, among other things, *being loose*, *being high-waisted*, and *being unflattering*. For ‘dad joke’/‘joke told by a dad’, they include properties like *being trite*, *being punny*, and *being groan-inducing*.⁹² They also, as the explicit etymological connections suggest, include stereotypes more basically associated with ‘mom’ and ‘jeans’ and ‘dad’ and ‘joke’, respectively—for example, *being uncool/embarassing*, *being made of denim*, *being silly*, and *being intended to be funny*. Familiarity with these stereotypes is absolutely essential for understanding exchanges like (86) and (87):

- (86) A: Obama wore mom jeans when he threw out the first pitch at a White

88. Twitter, 30 Mar 2012 [direct] [archive]

89. Twitter, 15 Jul 2014. [direct] [archive]

90. Twitter, 2 Apr 2019. [direct] [archive]

91. Twitter, 15 Oct 2019. [direct] [archive]

92. For an analysis of these stereotypes, see “[The Dubious Art of the Dad Joke](#)” (Luu 2019).

Sox game.

B: But Obama's not a mom!

A: Whatever, he still wore mom jeans. Threw like a girl too.⁹³

(87) A: Ellen's telling a lot of dad jokes.

B: But Ellen's not a dad!

A: Whatever, she's still telling dad jokes. Someone tell this woman she's not funny.⁹⁴

These, like the other DC-corrections we've seen, are felicitous, cooperative exchanges. And like in those other cases, this marks a contrast with the corresponding NPA-corrections:

(88) A: Obama wore jeans belonging to a mom when he threw out the first pitch.

B: But those were his jeans! And Obama's not a mom!

A: #Whatever, he was still wearing jeans belonging to a mom.

(89) A: Ellen's telling a lot of jokes [being] told by a dad.

B: But Ellen's not a dad!

A: #Whatever, she's still telling jokes [being] told by a dad.

And once again, this asymmetry is plausibly explained in terms of terms of *E*- and *S*-overlap.

Such evidence of systematicity, I have claimed, is *prima facie* evidence in favor of GENERALITY:

generality all intuitive DC/NPA pairs in English are unified by a general linguistic relationship.

But we might worry that (86)–(89) are not actually analogous to the other DC- and NPA-corrections we've considered so far. The pragmatic story I proposed for those earlier exchanges appealed to a default presumption among competent users. But we might question whether, in (86) and (87), B's replies are actually things that a competent user would (in earnest) say.

It is true, of course, that B's replies in (87) and (88) are, like B's replies in the analogue exchanges, *conversationally relevant*; it does not feel *random* or *inexplicable* to us, as competent speakers, that B replies the ways that he does. Still, it

93. Adapted from: Twitter, 3 Dec 2022. [[direct](#)] [[archive](#)]. For more on this notorious presidential scandal, See "[Obama on 'mom' jeans: 'I've been unfairly maligned'](#)" (Kim 2014).

94. Adapted from: Twitter, 13 Feb 2020. [[archive](#)]


is significantly less natural sounding, funny even.⁹⁵ Would a fully competent user of ‘mom jeans’ really be confused, even if only for a second, by the fact that Obama isn’t a mom?

I’m myself inclined to think no, actually. But happily, this poses no problem at all for DC/NPA OVERLAP. Recall that the basic linguistic relationship between DCs and NPAs countenanced by that hypothesis involves “substantial” degrees of *E*- and *S*-overlap:


DC/NPA overlap for any DC and NPA intuitively linked in meaning, the linguistic relationship between them involves substantial overlap in

- (a) semantic extension (*E*-overlap), and
- (b) the stereotypes, beliefs, and evaluative attitudes associated with each expression in the minds of competent speakers (*S*-overlap)

Initially, I illustrated the idea of *E*-overlap with a Venn diagram wherein only the outermost parts of the circles were not overlapping:



But for any given DC/NPA pair there are many other Venn diagrams we might draw, corresponding to equally as many possible instantiations of “substantial” *E*-overlap that pair might exhibit:



DC/NPA pairs, in other words, need not all exhibit exactly the same ratios of *E*- and *S*-overlap as ‘chick flick’/‘romantic comedy’ for DC/NPA OVERLAP to be true. It is consistent with that hypothesis that some DC/NPA pairs exhibit more *E*- and *S*-overlap than ‘chick flick’/‘romantic comedy’—even, perhaps, to the point of perfect co-extension—and that others exhibit less.

Hence, “exceptional cases”—occasions on which competent users take one member of a DC/NPA pair to correctly apply but not the other—may be more

95. Indeed, in the case of ‘dad joke’, it is actually the basis of a popular, well, dad joke:

Question: What do you call someone who tells dad jokes but isn’t a dad? Answer: A faux pa!

or less exceptional. Accordingly, it will be more or less conversationally strange for a speaker to presume such “exceptions” are to be bracketed by default. We may thus predict that speakers who are more competent with a given DC/NPA pair will have a stronger grasp on all of the following than speakers who are less competent:

- (a) the relevant associated stereotypes and attitudes;
- (b) the degree of *E*-overlap judged by other competent users to obtain between the relevant expressions; and
- (c) the norms and use practices of such speakers in ordinary conversation.

It is an intuitive datum that competence with DCs comes in degrees; DC/NPA OVERLAP can straightforwardly explain this.

5.3. Charge: *Difference in Object/Target Kind*

So far, all of the pairs examined in this section have involved group or category expressions for nonhuman objects. ‘Chick flick’ is a term for movies, not people. So why think its relationship to ‘romantic comedy’—whatever that relationship involves—should have any bearing on our thinking about paradigmatic slurs?

The same patterns I have been focusing on extend to DC/NPA terms for human persons or groups. Consider the DC/NPA pair ‘d*ke’/‘lesbian’. It is a prediction of our toy version of Williamson’s conventional implicature theory (and presumably his actual theory, as he intends it), that the semantic relationship between ‘d*ke’ and ‘lesbian’ involves strict co-extension. DC/NPA OVERLAP allows that this might be the case—but is only committed to the relevant *E*-overlap being substantial. And indeed, this latter, weaker prediction seems to accord much more comfortably with the actual practices of competent users.

Take, for example, a sentence like (49):

- (49) I always wanted a lesbian friend, lesbian not a dyke.

Anyone familiar with the term ‘d*ke’ as used a DC has almost certainly heard it used the way it is here—namely, to carve a distinction between the “acceptable” lesbians, and the unacceptable ones.

We may even be tempted to suppose for this reason that, contrary to both DC/NPA OVERLAP and the Williamsonian co-extension thesis, ‘d*ke’ actually picks out a proper subset of the individuals in the extension of ‘lesbian’ (cf. Ashwell 2016). But this would be to overlook the rest of the relevant data. It would be to overlook the conservative father who, upon learning that his daughter is bisexual,

calls her 'd*ke' in a rage. And it would be to overlook the casual, commonplace bigotry of utterances like (50):

- (50) Not all dykes are lesbians. I got a cousin who's a dyke but she has a husband.

Even more to our purposes here, it would be to overlook exchanges like (90):

- (90) A: Our daughter's a dyke.
 B: But she's bisexual, not a lesbian!
 A: Whatever, she's still a dyke.

Such exchanges, structurally, are just like the other DC-corrections we have seen; and just like those other DC-corrections, they are natural-sounding and immediately comprehensible. Speaker A in (91) is not merely believable as a competent user of 'd*ke', but is behaving exactly as those of us who know such speakers have come to expect—it makes no difference to him whether his daughter likes only girls, or not.

Finally, the (regrettable) normalcy of exchanges like (90) again stands in predictable contrast with (91), the corresponding NPA-correction:

- (91) A: Our daughter's a lesbian; she doesn't like boys.
 B: But she's bisexual, not a lesbian! She likes boys too!
 A: #Whatever, she's still a lesbian; she doesn't like boys.

And again, this asymmetry in felicity is not puzzling at all if our working hypothesis is DC/NPA OVERLAP. It *is* rather puzzling, however, if our working hypothesis is something like our generalized toy version of Williamson's co-extension thesis, or the proper subset view (hereafter NPA SUBSET):

DC/NPA coextension for any DC and NPA intuitively linked in meaning, the linguistic relationship between them involves semantic coextension.

NPA subset for any DC and NPA intuitively linked in meaning, the extension of the DC is a *proper subset* of the semantic extension of the NPA.

DC/NPA COEXTENSION predicts that 'd*ke' and 'lesbian' are extensionally equivalent in literal uses—that is, that all and only lesbians are d*kes—and that competent users know this; NPA SUBSET predicts that all literal referents of 'd*ke' are lesbians—but not all lesbians are literal referents of 'd*ke'—and that competent users know this. So for either hypothesis to be correct, speakers like the

conservative father in (91) must be confused about the meaning of ‘d*ke’, or using it in a nonliteral way. We should reject the former out of hand. Speakers like the conservative father, and utterances like (49) and (50), are not strange or idiosyncratic; as I have taken some pains in this paper to show, they are utterly banal. To dismiss such cases as confused or semantically defective is to posit *rampant* linguistic incompetence within the very speech communities supposedly at issue—*viz.*, communities wherein the target DCs are actually, routinely used. Methodologically, this is a nonstarter.

The latter proposal, that utterances like (49) and (50) and exchanges like (91), involve *nonliteral* uses of the relevant DCs, is more serious. Still, I don’t think we have reason to prefer it over an overlap-theoretic alternative.

6. Why Not Say It’s Metaphor?

Probably no one actually accepts DC/NPA COEXTENSION in its full generality—but many extant theories do entail that ‘lesbian’ and ‘d*ke’ are truth-conditionally equivalent in literal uses. And if one arrives with such prior theoretical commitments, a natural response to the data I have presented is to ask: why not say it’s metaphor? Perhaps uses I have been focusing on are simply nonliteral. Indeed, if we’re antecedently committed to something like Williamson’s conventional implicature view, then we *have* to say such data involves nonliteral uses. Sentences like (92) and (93) cannot, if ‘d*ke’ and ‘lesbian’ are extensionally equivalent, be literally true:

(92) She’s a dyke but not a lesbian.

(93) She’s a lesbian but not a dyke

So, if we think something like Williamson’s view is true for DC/NPA pairs like ‘d*ke’/‘lesbian’, then we have reason to think uses like those in (93) and (94) are *contracted* and *extended*, respectively.⁹⁶

But do we have *independent* reason to think this?

As a first pass test, I offer that competent speakers generally *grasp* when they are speaking metaphorically. Indeed, it seems we pronounce a metaphor “dead” precisely when and because its “nonliterality,” however retrievable it may still be in principle, has now been so fully lexicalized away as to go virtually unnoticed in practice (e.g., ‘mouth of the river’). “Live” metaphors, by contrast, have a distinct “figurative feel.” For example, consider (94):

96. Jeshion 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 (2013b: 251–53).

- (94) 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.⁹⁷

As competent speakers of English, we *know* by the end of this exchange that A intends a nonliteral meaning of 'girl'. Moreover, we have an intuitive sense of the *way* that B is failing to understand when he attempts to correct A—viz., by taking as literal an utterance which (in retrospect) was clearly supposed to be figurative. It is simply *obvious* to us, by the end of the exchange, that a literal reading of 'Jack's a girl' is not available—even if it might have been initially.

Indeed, as hearers, we may reasonably wonder at the *beginning* of (94) about the intended literality of A's initial claim. But once we imagine that *we* are A, there is simply *no question* of how 'girl' is being used. To say that that Jack is (literally) a girl and to say that Jack is (figuratively) a girl are two very different things—whether we are saying one or the other is something that *we*, as competent speakers, would know.

So if the uses of DCs in sentences like (49) and (50) and exchanges like (90) are in some sense metaphorical (either extended or contracted), then *either* (a) they are metaphorical in the sense that "dead" metaphors are metaphorical, or (b) they are "live" metaphors.

- (49) I always wanted a lesbian friend, lesbian not a dyke.
 (50) Not all dykes are lesbians. I got a cousin who's a dyke but she has a husband.
 (90) A: Our daughter's a dyke.
 B: But she's bisexual, not a lesbian!
 A: Whatever, she's still a dyke.

If (a), then DCs in sentences like (49) and (50) and exchanges like (91), like 'mouth of the river', are (now) fully-lexical terms in their own right, with a (literal) semantic relationship to NPAs which is (still) *prima facie* best explained by an overlap thesis. If (b), then DCs in sentences like (49) and (50) and exchanges like (90) *aren't* fully-lexical in their own right, and the speakers behind those DCs generally realize—in virtue of being competent users, as again we ought to assume—that they're speaking metaphorically.

Whether views like Williamson's are more plausible than overlap views turns on whether this latter prediction about DC-users', and in particular slur-users', knowledge and intentions is true. And for what it's worth, as someone

97. This is a modified version of an example from Jeshion (2013b).

who grew up around such speakers, I do not think that it is. As someone raised around slur users, I *think* I know what they mean when they say things like (49) and (50), or (51) and (52), or when they use the n-word.

(51) There are gays and then there are faggots.⁹⁸

(52) Whoever decided to allow transgender faggots in the military is a retard.⁹⁹

The supposition that such speakers—like A in (95) with ‘girl’—are *knowingly* and *intentionally* using the relevant slurs in nonliteral ways strikes me, frankly, as wildly implausible. They are simply using them, as they do, in the way that they take them to mean.

That so many theorists may be tempted to think otherwise, is, I suspect, the result of their taking a very *extreme* kind of DC-user as their paradigm. This is the avowed white supremacist, the self-righteous and evangelizing zealot, whose unvarnished hatred for a certain perceived “kind” of fellow human being is not only conscious and endorsed, but *directed* and *absolute*. But most bigotry is not so unyielding. On the contrary, it is the *business* of ordinary bigotry to “make exceptions”—to distinguish “model minorities” from “the riffraff,” to separate “the good ones” from “the bad.” Its willingness to arbitrate, often in one-off, confabulatory, and inconsistent ways, who should and should not “count” as a genuine target of contempt is not a secondary or peripheral feature, but a *central hallmark*. Everyday bigotry, and everyday slur-use, is prototypically *slippery* and *capricious*. This is what makes it so insidious, and thus should be at the center of our academic theorizing. An overlap view vindicates this centering in a way that more orthodox views like Williamson’s do not.

7. Additional Objections and Replies

Appealing to non-literality to explain the target data thus comes at significant intuitive and methodological cost. There are other ways the one might try to explain the data—though these too, I suggest, are less plausible than an overlap account.

For example, perhaps sentences of the kind I have focused on are like (95):

(95) He’s a novelist, but he’s not a *novelist*.

98. Twitter, 17 Nov 2017. [[direct](#)] [[archive](#)]

99. Twitter, 30 Jun 2016. [[direct](#)] [[archive](#)]

Such sentences seem, semantically, like they should be contradictory, and therefore infelicitous. But they are not infelicitous, because they are pragmatically rescued—what (96) really means is something like (96):

(96) He's a novelist, but he's not a *good/characteristic* novelist.¹⁰⁰

The pragmatic mechanism involved here is slightly different from the mechanism involved in metaphor, and turns on the presence of intonation or stress. Indeed, for (95) to be felicitous, there *must* be a difference in stress between the two instances of 'novelist'; otherwise, the sentence is unintelligible:

(97) #He's a novelist, but he's not a novelist.

This need for a difference in stress is explained by—and seemingly triggered by—the *fact* that 'novelist' and 'novelist' mean the same thing. So if sentences like (93) are akin to (95), we should expect an analogous difference in stress to be required for felicity:

(93) She's a lesbian but not a dyke.

Of course, someone *might* utter a version of (93) with stress on the last word. But this difference in stress is not *required* for (94) to be felicitous, as it is for (95), and is often absent as a matter of fact.

This is one reason to doubt that the target data can be explained along the lines of (95). Another is that such an explanation can only account for *one direction* of DC/NPA divergence. While it might be able to account for sentences like (94)—so that (94) means something like (98)—it cannot account for sentences like (93):

(94) She's a lesbian but not a dyke.

(98) She's a lesbian but not a *bad/stereotypical* lesbian.

(93) She's a dyke but not a lesbian.

To explain (93), one would need to appeal to a different mechanism (presumably metaphor), thus multiplying theoretical commitments and linguistic machinery to account for data that an overlap thesis can explain outright.

100. Thanks to an anonymous referee for raising this objection.

Finally, assimilating sentences like (94) to (95) leads to false predictions about how willing DC-users will be to admit that something belongs to the relevant category.

- (94) She's a lesbian but not a dyke.
 (95) He's a novelist, but he's not a *novelist*.

If pressed, the speaker of (96) will (perhaps begrudgingly) admit that the novelist in question is, in fact, a novelist—just not a very good one. The same is not plausible about (94) or the rest of the target data. Speakers of (1) and (2) are not similarly obligated, even under pressure, to admit that the 2016 *Ghostbusters* is a romantic comedy and *Silver Linings Playbook* is a chick flick, respectively:

- (1) The new *Ghostbusters* is chick flick, but it isn't a romantic comedy.
 (2) *Silver Linings Playbook* is a romantic comedy, but it isn't a chick flick.

This is clearest if we imagine situations, like bets, where the stakes are fairly high for whether something (actually) belongs to a category or not. *Contra* Elisabeth Camp (2013), who appeals to bets as evidence for something like DC/NPA COEXTENSION, DC-users are liable to have *genuine disputes* about DC category boundaries in betting contexts. To suppose that bets about how many so-and-sos are in a building will be settled quickly and easily between bigots on the basis of something like DC/NPA COEXTENSION is to assume something naive, and false, about everyday bigotry. Such bets are liable to admit of the same kind of disagreement as bets between 'chick flick'-users about how many "chick flicks" there are. If A and B have a bet about how many "chick flicks" are in B's DVD collection, A and B might *easily* disagree about who has won—and not simply because one is being unreasonable or obtuse. The assumption that DCs and NPAs are extensionally equivalent leads to badly wrong predictions about possible bets made in terms of those expressions, by predicting that such bets would be easily and uncontroversially discharged as won or lost by competent users.

Tractable, winnable bets require clear winning and losing conditions; hence, making a bet puts pressure on us to use well-defined terms. A bet that turns on whether a given DC applies will in general not be well-defined. We should expect such bets not only to be more *contentiousness* among actual DC-users than Camp and other theorists sympathetic to DC/NPA COEXTENSION seem to assume, but also less *natural* for DC-users to make in the first place. If we would not agree to a bet about how many "chick flicks" are in our friend's DVD collection without first getting clear on what the relevant criteria are, we should not expect users of more paradigmatic slurs to agree to analogous bets without first establishing, in more precise terms, what it is, exactly, that they mean.

8. What It Means to Accept Overlap

8.1. Taking Stock

It will be helpful, at this point, to review the dialectic. I began by observing, first, that for competent speakers of English, there is a wide range of group or category expressions with recognizably pejorative uses (DCs); and second, that among these group or category expressions, many are such that competent speakers associate them intuitively with other, more “neutral” group or category expressions (NPAs). I then provided a number of examples of such pairs of expressions, including paradigmatic slur/“neutral counterpart” pairs, as well as a wide selection of data sentences in which the constituent terms are used in explicitly contrastive ways.

Given the *prima facie* systematicity of such data, and naturalness of what I called **OVERLAP HYPOTHESIS** for ‘chick flick’/ ‘romantic comedy’, I suggested the following generalized version as an “obvious first thought” for *all* intuitively-given DC/NPA pairs:

DC/NPA overlap for any DC and NPA intuitively linked in meaning, the linguistic relationship between them involves substantial overlap in

- (a) semantic extension (*E*-overlap), and
- (b) the stereotypes, beliefs, and evaluative attitudes associated with each expression in the minds of competent speakers (*S*-overlap).

I then argued that this hypothesis is more plausible, because more parsimonious and more faithful to the actual practices of everyday speakers, than two rival general hypotheses, **DC/NPA COEXTENSION** and **NPA SUBSET**:

DC/NPA coextension for any DC and NPA intuitively linked in meaning, the linguistic relationship between them involves semantic coextension.

NPA subset for any DC and NPA intuitively linked in meaning, the extension of the DC is a *proper subset* of the semantic extension of the NPA.

These rival hypotheses, I argued, require additional theoretical commitments and machinery to explain the same data that **DC/NPA OVERLAP** can explain with only basic Stalnakerian assumptions about common ground. Moreover, I argued that the additional required commitments are at best non-trivial, and at worst unmotivated.

Of course, this is not to say that **DC/NPA COEXTENSION** and **NPA SUBSET**—or the more circumscribed versions thereof actually defended by theorists—are deci-

sively *false*; to be very clear, I do not claim to have shown this. What I do claim to have shown is that, in the absence of defeating evidence, we have significantly stronger *prima facie* reason to accept DC/NPA OVERLAP over either narrower alternative.

Now, supposing you are willing to take all of that on board—what, exactly, does “accepting” DC/NPA OVERLAP commit us to? What does accepting an “overlap view” about DC/NPA pairs actually involve?

A complete semantic theory of DC/NPA pairs is beyond the scope of this paper; and to be very clear on this point as well, I do not take myself to have given one. That being said, there are some clear constraints that accepting DC/NPA OVERLAP places on the relevant space of semantic views. I’ll do two things to clarify the options.

First, I will sketch some semantic frameworks amenable to DC/NPA OVERLAP. As it turns out, however, even these options are not the only paths forward from the arguments in this paper. I hinted before that DC/NPA OVERLAP was not the “overlap thesis” I ultimately wished to defend—we are now in a position to see why. Despite being the “obvious first thought,” DC/NPA OVERLAP it is actually stronger in its semantic commitments than is strictly necessary to explain the data. So, having first sketched some ways to accommodate those commitments, the second thing I will do is identify a less “semantically-loaded” alternative.

8.2. Making Semantic Sense of E- and S-Overlap

DC/NPA OVERLAP, recall, hypothesizes that the linguistic relationship between any given DC and NPA intuitively linked in meaning involves a substantial degree of *E*-overlap—viz., overlap in *semantic extension*—as well as a substantial degree of *S*-overlap—viz., overlap in *associated stereotypes, beliefs, and attitudes*.

To my knowledge, no one in the philosophical literature has previously argued for DC/NPA OVERLAP in its fully general form. Adam Croom (2011; 2015), however, has defended something very *close* to DC/NPA OVERLAP for paradigmatic slurs and so-called “neutral counterpart” terms—and for very similar reasons; namely, data like (51) and (52):

(51) There are gays and then there are faggots.¹⁰¹

(52) Whoever decided to allow transgender faggots in the military is a retard.¹⁰²

101. Twitter, 17 Nov 2017. [\[direct\]](#) [\[archive\]](#)

102. Twitter, 30 Jun 2016. [\[direct\]](#) [\[archive\]](#)

Such instances of what I called “bidirectional divergence” in ascriptions by competent users have motivated Croom and more recently Eleonore Neufeld (2019) to develop *prototype theories* of the meanings of paradigmatic slurs. Though not the only option, prototype semantics is a natural fit for DC/NPA OVERLAP, and Croom’s proposal in particular seems easily adjusted to accommodate it. (As we’ll see, Neufeld’s proposal is consistent with a more minimal version.)

Prototype semantic theories are motivated by empirical work on concepts and categorization and embrace the idea of “fuzzy”, or borderline, category membership (Rosch 1978; Rosch & Mervis 1975; Coleman & Kay 1977). Such theories reject classical “dictionary definition” frameworks according to which category membership is determined by jointly necessary and sufficient application conditions, instead holding that lexical entries encode highly structured prototype concepts governed by similarity relations.¹⁰³ Accordingly, Croom (2011; 2015) has proposed a “family resemblance view” on which slurs encode a “constellation” of ranked prototypical features (or stereotypes), satisfaction of any number of which may be contextually sufficient, but is crucially *not* necessary, for category membership (2015: 34). Thus, the view can explain competent users’ judgments that sentences like (51) and (52) are (or in principle could be) true, and why *E*-overlap between the relevant expressions is substantial but not complete: the mere fact that someone counts as a member of the relevant “neutral counterpart” category does not mean they count as a member of (or are in the extension of) the associated slur category (and vice versa). And while Croom is not explicit about what the corresponding semantics looks like for the relevant “neutral counterpart” terms, it is clear enough in that overlap in associated stereotypes (what I’ve called *S*-overlap) is supposed to be what grounds the relevant semantic link to the corresponding slurs. It is thus easy to see how his or a similar prototype account could be developed into an account of DC/NPA OVERLAP.

Does that mean that accepting DC/NPA OVERLAP requires explaining *S*-overlap *semantically*? Not necessarily! Prototype views are natural candidates for explaining the role of *S*-overlap in the literal meanings of DCs and NPAs, because they make the relevant overlapping stereotypes *part* of the literal meanings. But as I have described *E*- and *S*-overlap, DC/NPA OVERLAP is compatible with a range of semantic frameworks which are *not* predicated on (controversial) prototype theory, and may account for the role *S*-overlap differently.

Inferential or “*relevance-theoretic*” frameworks, for example, may account for *S*-overlap wholly non-semantically, by locating the relevant beliefs, stereotypes, and representations in “mental files” attached to atomic lexical concepts (Sperber

103. For an attempt to model the relevant similarity relation, see Decock and Douven (2014).

& Wilson 1995; Wilson & Sperber 2006; Wilson & Carston 2007).¹⁰⁴ The atomic concepts themselves, understood purely denotatively, would in turn be sufficient to explain *E*-overlap, while continuing to feed straightforwardly into more traditional truth-conditional evaluation.

Attempting to thread the needle between prototype theory and traditional truth-conditional theories is another option still, *dual-content semantics*. Dual-content frameworks posit *both* a denotative structure that atomically determines extension (an “E-structure”), *and* a prototype or conceptual structure of associated beliefs or representations (a “C-structure”) (Del Pinal 2015; Del Pinal 2018; Pustejovsky 1995). The potential for mapping of E- and C-structures to what I’ve called *E*-overlap and *S*-overlap for DC/NPA pairs is straightforward.¹⁰⁵

Finally, and more generally, an overlap thesis for DC/NPA pairs is compatible with different conceptions of application conditions, which may or may not amount to conditions for “truth.” An assertability semantics, for example, could vindicate the linguistic *competence* of DC-users, even if we think no DC-ascription is, or ever could be, actually true.

8.3. An Even More Minimal Alternative

Indeed, when it comes to paradigmatic slurs at least, the idea that they are *extensionally empty* is one that philosophers often find attractive. This idea has been most prominently defended by Christopher Hom (2008; Hom & May 2018), but its tension with DC/NPA OVERLAP is brought out most sharply by Eleonore Neufeld’s (2019) recent prototype proposal, which by all other lights would seem like a natural candidate for being developed into an “overlap view.”

Like Croom (2011; 2015), Neufeld takes examples like (51) and (52) to be central (literal) use cases of paradigmatic slurs and so-called “neutral counterpart” terms. And like Croom, she has offered a prototype semantics on which the relevant expressions encode “family resemblance” concepts structured around stereotypes. On her view, though, paradigmatic slurs are *unlike* their corresponding “neutral counterpart” terms in that, characteristically, they encode “mini theories” which posit *intrinsic essences* causally linked to the relevant (negatively-valenced) stereotypes. And since no person actually has such an essence, paradigmatic slurs turn out to never actually refer. As lexical concepts, their “fuzzy” membership boundaries can vindicate the literality of ascriptions like (51) and (52), as well as competent users’ *beliefs* that such ascriptions are true. But the

104. A broadly inferentialist view of (paradigmatic) slurs has been given by Tirrell (1999).

105. Zeman (2021) has recently defended a different application of dual-content semantics to slurs on which they are highly polysemous. Zeman’s focus is on variation in meaning among uses of the *same* slur, rather than the relationship between slurs and so-called “neutral counterparts.”

ascriptions themselves, and the “mini theories” they presuppose, will always turn out to be false.

It is actually not totally clear what the corresponding proposal for “neutral counterpart” terms is supposed to be for Neufeld; but assuming that they are not also ultimately non-referring,¹⁰⁶ it seems clear that her view is in tension with DC/NPA OVERLAP, which posits *actual* extensional overlap between DCs and (*ex hypothesi*) non-empty NPAs:

DC/NPA overlap for any DC and NPA intuitively linked in meaning, the linguistic relationship between them involves substantial overlap in

- (a) semantic extension (*E*-overlap), and
- (b) the stereotypes, beliefs, and evaluative attitudes associated with each expression in the minds of competent speakers (*S*-overlap).


So if we are attracted to something like Neufeld’s view, or the more general thought that paradigmatic slurs (at least) are extensionally empty, then the various “paths forward” from this paper may seem like just as many dead ends.

106. Neufeld is very clear in stating that slurs and their so-called “neutral counterparts” are *not synonymous* (2019: 5); and that, if neutral counterpart terms are also essentializing, that they do not essentialize in the same way slurs do (2019: 6). It is not clear, however, if this means that neutral counterpart terms simply encode a *different* essentialist “mini theory” than their corresponding slurs, or not. If they do, and it turns out that “mini theory” is *also* false, then Neufeld’s account could actually be (trivially) consistent with DC/NPA OVERLAP. Indeed, to the extent that Neufeld’s proposed semantics for slurs is plausible, and to the extent that the cognitive science research on essentialization she draws on actually supports it, an analogue proposal on which so-called “neutral counterparts” also encode a falsely essentializing “mini-theory” seems like it should be equally well-supported. At the very least, it is unclear why evidence of (especially) racial essentialization in cognition should support an essentialist prototype theory of *slurs*, but not “neutral counterparts,” when speakers who use both are in general *not* “neutral” about either. Neufeld says something suggestive along these lines when she writes:


[E]ven if we say that slurs’ neutral counterparts are sometimes essentialized, I contend that the essence referred to by, say, ‘chink’ is not identical to the essence of ‘Chinese’. Also under the assumption that ethnic labels sometimes function as kind terms, the kind they designate differs from the kind their closest slur-relative designates. I merely chose ‘Chinese essence’ as a label for the essence placeholder that unifies, *in the eyes of the racist*, the alleged referents of the slur. (2019: 5; my emphasis)

There seems room, that is, to think that, in addition to postulating a “ch*nk essence,” the speaker who uses that slur also postulates a (different, but similar) “Chinese essence” that (in their eyes) unifies the referents of that term. But this would seem to imply that slur users are either *incompetent* with the “actual” (non-racist, non-essence-positing) meanings of “neutral counterpart” terms, or else *speak falsely* when using “neutral counterpart” terms for (structurally) the same reason they speak falsely when using slurs. And while this version of the view could be, at least in principle, compatible with DC/NPA OVERLAP, it does not strike me as especially promising.


In fact, however, the central arguments in this paper don't turn on DC/NPA OVERLAP, and accepting them does not require positing actual *E*-overlap. I have focused on DC/NPA OVERLAP until this point because its contrast with DC/NPA COEXTENSION and NPA SUBSET is direct; and in the context where these two alternatives are salient, it is the "obvious first thought" to have about many of the examples I have considered:



NPA subset. All chick flicks are romantic comedies, but not all romantic comedies are chick flicks.



DC/NPA overlap. In general, most chick flicks are romantic comedies, and most romantic comedies are chick flicks.



DC/NPA coextension. All chick flicks are romantic comedies, and all chick flicks are romantic comedies.

And while I maintain that DC/NPA OVERLAP is *better able* than these alternatives to explain the relevant data, it is also *strictly stronger* than is needed to explain that data.

Recall that, in the case of 'chick flick'/'romantic comedy', this data included the following:

- (a) that competent users of 'chick flick' and 'romantic comedy' take those expressions to be tightly intuitively related, both in extension and in the stereotypes and attitudes associated with them;
 - (b) that, despite this fact, competent users sometimes apply 'chick flick' and 'romantic comedy' to different things;
 - (c) that exchanges like (76) are, to competent users of 'chick flick' and 'romantic comedy', familiar and immediately comprehensible;
- (76) A: The new *Ghostbusters* is a chick flick.
 B: But it's an action movie, not a romantic comedy!
 A: Whatever, it's still a chick flick.
- (d) that B's reply in (76) seems to presuppose that all and only chick flicks are romantic comedies (viz., that *E*-overlap between 'chick flick' and 'romantic comedy' is complete)
 - (e) that despite this fact, A's rejoinder in (76) is not only felicitous, but is consistent with, and indeed actually suggestive of, basic competence with the relevant expressions; and

- (f) that this marks a clear asymmetry with (79), where the exchange in general is defective, and A's rejoinder in particular is bizarre and uncooperative.
- (79) A: The new *Ghostbusters* is a romantic comedy.
 B: But it's an action movie, not a romantic comedy!
 A: #Whatever, it's still a romantic comedy.

I argued that DC/NPA OVERLAP together with some basic assumptions about common ground, easily explains *all* of (a)–(f) by explaining how competent users are likely, in general, to observe a default but defeasible conversational presumption of *complete E-overlap* between 'chick flick' and 'romantic comedy' in practice, while in principle being open to possible exceptions.

Notice, however, that to give this explanation, *we need not posit any actual E-overlap at all*. All that is required is that *competent users* believe there is *E-overlap*. Like Neufeld, we could think this belief is false (and if not about 'chick flick'/'romantic comedy', then perhaps about other DC/NPA pairs). After all, it would be strange if it were *impossible* for a DC like 'witch' to have an intuitive NPA, strictly because its extension has turned out to be empty. So all we need to explain (a)–(f) is a *presupposition* of substantial *E-overlap* among competent users, along with some adequate theory of *S-overlap* in which to ground it. Generalizing to other DC/NPA pairs, call this more minimal hypothesis PRESUPPOSED DC/NPA OVERLAP (hereafter PRESUPPOSED OVERLAP):

presupposed DC/NPA overlap for any DC and NPA intuitively linked in meaning, the linguistic relationship between them is grounded in

- (a) a (possibly false) presupposition among competent users that the DC and NPA substantially" overlap in semantic extension

where this is in turn grounded in

- (b) a substantial degree of (actual) overlap in the stereotypes, beliefs, and evaluative attitudes these users associate with each expression.

The central difference between DC/NPA OVERLAP and PRESUPPOSED OVERLAP is that PRESUPPOSED OVERLAP carries no particular commitments about which (if any) ascriptions of DCs are true. It could turn out that speakers are simply wrong in their beliefs about the extension of a given DC (or indeed, even of a given NPA). It could turn out that some, many, or perhaps even all DCs are like 'witch' with respect to their (actual) truth conditions. If anyone is competent with DCs, DC-using speech communities are. Linguistic competence among DC-users need not involve actually latching onto the extensional facts.

9. Overlap vs. Orthodoxy

PRESUPPOSED OVERLAP is thus a simple, elegant, *flexible* hypothesis about DC/NPA pairs that captures a huge range of everyday linguistic data. My main contention in this paper is that, when it comes to the more specific case of paradigmatic slurs, it is also *the most we should assume* about their semantics, and in particular their relationship to so-called “neutral counterparts.” And this thesis, it turns out, is surprisingly heterodox.

Despite their many points of disagreement, nearly all existing theories of (paradigmatic) slurs take something surprisingly substantive for granted about the relationship words like ‘k*ke’ bear to words like ‘Jew’ or ‘Jewish’; namely, that this is a *semantically distinctive* relationship, which at once *unifies* paradigmatic slurs as a *special semantic class* while (in effect) carving them semantically *away* from other DCs with intuitive NPAs. This “special” relationship is cashed out in different ways by different theorists. But in every case, it is assumed that understanding this relationship for *one* slur/“neutral counterpart” pair is *directly importable* to other such pairs, such to know the (semantic) meaning of a given “neutral counterpart” term is *ipso facto* to know the (semantic) meaning of the corresponding slur. No additional familiarity with stereotypes is required.¹⁰⁷

This assumption is so diffuse, and so deeply backgrounded, in current philosophical theorising that it is difficult to articulate precisely. As I have already suggested, it is clearest in the case of views like Williamson’s which, if extended to accommodate GENERALITY, would entail DC/NPA COEXTENSION:

generality all intuitive DC/NPA pairs in English are unified by a general linguistic relationship.

DC/NPA coextension for any DC and NPA intuitively linked in meaning, the linguistic relationship between them involves semantic coextension.

According to these views, the semantic (truth-conditional) content of a slur *just is* the semantic (truth-conditional) content of its so-called “neutral counterpart”; 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.¹⁰⁸ Following Croom (2015), I will call this thesis *coreferential-*

¹⁰⁷ Some early theories (most notably Hom 2008) built stereotypes into the meanings of slurs; however, since Jeshion (2013b)’s influential critique, stereotype views have fallen out of favor.

¹⁰⁸ These include: (a) hybrid expressivist views, according to which the meanings of slurs involve the descriptive meanings of their NPA plus a negative attitude (Saka 2007; Richard 2008; Jeshion 2013a) or gestural content (Hornsby 2001); (b) implicature views, according to which slurs conventionally imply something negative that their NPAs do not (e.g., Williamson 2009; McCready

ism.¹⁰⁹ Coreferentialist views, then, assume there is a *simple recipe* for computing the (truth-conditional) meaning of a given slur from the (truth-conditional) meaning of the relevant “neutral counterpart” term. From there, all that is required for competence with the slur is to recognize that it, but *not* its “neutral counterpart”, is derogatory in typical uses.¹¹⁰

Coreferentialism is the dominant position in the philosophical literature. But even theorists who have proposed non-coreferentialist views have not questioned the assumption of a general “recipe” for computing the meanings of (derogatory) slurs from the meanings of their so-called “neutral” (non-derogatory) counterparts; they have only questioned the ingredients list. For example, according to Kent Bach’s *loaded descriptivism*, the meaning of a slur, *s*, is the meaning of its neutral counterpart, *n*, plus a “side comment” along the lines of “[*n*]s are contemptible in virtue of being [*n*]” (2018: 64).¹¹¹ Christopher Hom and Robert May (2013; 2018) propose 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. Ashwell (2016) suggests a version of NPA SUBSET. On these views the details of the recipe are different, but it still involves (a) *fixing* the content of the slur via, or in terms of, the content of the relevant “neutral counterpart” term; and (b) *adding* a derogatory element. Call this the JUST-ADD-BAD assumption about the relationship between slurs and so-called “neutral counterpart” terms.

Getting a gestalt feel for JUST-ADD-BAD is a bit easier than making it technically precise. But if we wish to make it precise, it is something like the following:

just-add-bad Assume a well-defined set (call it *S/N*) of all and only slur/“neutral counterpart” pairs. Then, for any pair of expressions *s/n*, if

2010; Whiting 2013); (c) prohibition views, according to which the only difference between slurs and their NPAs is taboo (e.g., Anderson & Lepore 2013a; 2013b); 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” signalling endorsement of associated attitudes and ideologies (e.g., Bolinger 2017; Nunberg 2018).

¹⁰⁹ Cf. DiFranco (2015), who calls this Neutral Counterpart Theory. In their response to DiFranco, Caso and Lo Guercio (2016: 265–66) define the position as a conjunction of 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. As they observe, (C₁) is required if coreferentialist theories are to apply to all slurs.

¹¹⁰ How exactly to account for this difference in derogatory force is where coreferentialist views diverge—see footnote 106, above.

¹¹¹ Bach also explicitly leaves it open whether slur ascriptions are “true of whomever [relevant] neutral counterpart is true of” (2018: 61).

$s/n \in S/N$, 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 are 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 derogatory, is sufficient for competence with s .

This is somewhat cumbersome, as the scope of (otherwise heterogenous) views it is intended to capture is very wide. But despite their numerous and complex differences, all such views have, for my purposes here, one important thing in common: *they are facially false* if extended to a broader range of DC/NPA pairs.

Consider again some of the data presented in §4:

- (1) The new *Ghostbusters* is chick flick, but it isn't a romantic comedy.
- (2) *Silver Linings Playbook* is a romantic comedy, but it isn't a chick flick.
- (45) Being a dad with "actually funny jokes and not just dad jokes" might be the best feedback you can get from your kid.
- (46) I don't understand why Cory Booker tells so many Dad jokes when he's not actually a Dad.
- (39) The Dark Knight is the rare instance of a superhero movie that isn't capeshit.
- (40) It's a sad fact to come to terms with but Star Wars is "capeshit".
- (14) It's okay to be liberal, but it's not okay to be a libtard.
- (15) I'm OK with people who support Bernie. I do not like Bernie Bros.
- (35) Baby-boomer is the generation. Boomer (by itself) is an insult to connote a state of mind. You're a baby-boomer, but not a boomer.
- (36) Mayo Pete is a boomer at 37. Does this help everyone understand that boomer isn't just an age thing?
- (73) I want a roommate that smokes weed but isn't a stoner.¹¹²
- (74) You can be a stoner without smoking weed.¹¹³
- (43) Not all women are cunts. Term is reserved for only the deserving.
- (44) Men can be cunts too. Just as they can be pussies.

¹¹² Twitter, 31 Jan 2020. [[direct](#)] [[archive](#)]

¹¹³ Twitter, 7 Apr 2020. [[direct](#)] [[archive](#)]

I have argued that the range of these examples, together with their systematic behavior in DC- and NPA-corrections, gives us strong (if defeasible) *prima facie* reason to assume that GENERALITY is true. If we take GENERALITY seriously, however, then it is clear that neither coreferentialist nor non-coreferentialist views that assume JUST-ADD-BAD will be tenable as a fully general theory. It is, I submit, 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.” Likewise obviously false are the analogue claims about ‘capeshit (movie)/superhero movie’, ‘dad joke’/‘joke told by dads’, ‘d*ke’/‘lesbian’, etc. DC/NPA OVERLAP and PRESUPPOSED OVERLAP more generally are not bound to JUST-ADD-BAD, and thus can easily accommodate the full range of data in a straightforward, unified way that orthodox views of (paradigmatic) slurs and so-called “neutral counterparts” cannot.

10. Diagnosis and Conclusion

This leaves us with the question: if JUST-ADD-BAD is such an implausible assumption, 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.¹¹⁴

These are what Jeshion (2013b) aptly calls “weaponized” uses of slurs, or cases like (5):

- (5) [Shouted at a gay couple holding hands]:
You’re going to hell, faggots!

¹¹⁴ Nunberg 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.” (2018: 241)

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 bigots, and ordinary bigotry “makes exceptions.” Indeed, this is precisely Kate Manne’s (2017) 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 some men (especially gay and effeminate men) “are.” By distinguishing the “good” women from the “bad” ones, misogyny insulates itself from charges of absolute sexism while (re)tightening its insidious 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”.¹¹⁵ 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 ‘Bernie Bro’, ‘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 JUST-ADD-BAD. 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.

Acknowledgments

Versions of this paper were presented at Robin Jeshion’s “Slurs” seminar at University of Southern California, GAP.10, Cognitive Structures 2018 at Heinrich Heine University, The Mid-Atlantic Philosophy of Language (MAPL) Workshop at West Virginia University, Nikki Ernst’s Words Workshop, the USC Philosophy Speculative Society, and Mark Schroeder’s advisee workshop. I would like to thank the audiences, whose questions, suggestions, linguistic intuitions, and

¹¹⁵ That exception-granting is central to the social and epistemic perpetuation of bigotry is emphasized by Begby (2013).

proffered examples greatly shaped the paper over the years, as well as the various organizers. For their contributions to the present version, I am grateful to the editors of *Ergo*, two anonymous reviewers, and copyeditor Robert Mason. Finally, I would like to thank Dan Pallies, Liam Bright, Jonathan Ichikawa, Christa Peterson, Andrew Stewart, John Hawthorne, Steve Finley, Alexis Wellwood, Carol Foster, and Melissa Head for their continued feedback and encouragement. To Mark Schroeder and Robin Jeshion I am indebted most of all.

References

- Allen, Sophie, Elizabeth Finneron-Burns, Jane Clare Jones, Holly Lawford-Smith, Mary Leng, Rebecca Reilly-Cooper, and Rebecca Simpson (2018, August 27). Derogatory Language in Philosophy Journal Risks Increased Hostility and Diminished Discussion. *Daily Nous*. <https://dailynous.com/2018/08/27/derogatory-language-philosophy-journal-hostility-discussion/>
- Anderson, Luvell and Ernie Lepore (2013a). Slurring Words. *Noûs*, 47(1), 25–48.
- Anderson, Luvell and Ernie Lepore (2013b). What Did You Call Me? Slurs as Prohibited Words. *Analytic Philosophy*, 54(3), 350–63.
- Ashwell, Lauren (2016). Gendered Slurs. *Social Theory and Practice*, 42(2), 228–39.
- Bach, Kent (2018). Loaded Words. In David Sosa (Ed.), *Bad Words Philosophical Perspectives on Slurs* (60–80). Oxford University Press.
- Begby, Endre (2013). The Epistemology of Prejudice. *Thought: A Journal of Philosophy*, 2, 90–99.
- Bolinger, Renee (2017). The Pragmatics of Slurs. *Noûs*, 51(3), 439–62.
- Camp, Elisabeth (2013). Slurring Perspectives. *Analytic Philosophy*, 54(3), 330–49.
- Caso, Ramiro and Nicolás Lo Guercio (2016). What Bigots Do Say: A Reply to DiFranco. *Thought*, 5, 265–74.
- Coleman, Linda and Paul Kay (1977). Prototype Semantics: The English Word Lie. *Language*, 57(1), 26–44.
- Croom, Adam (2011). Slurs. *Language Sciences*, 33, 343–58.
- Croom, Adam (2015). The Semantics of Slurs: A Refutation of Coreferentialism. *Amper-sand*, 2, 30–38.
- Davis, Christopher and Elin McCready (2020). The Instability of Slurs. *Grazer Philosophische Studien*, 97(1), 63–85.
- Decock, Lieven and Igor Douven (2014). What Is Graded Membership? *Noûs*, 48(4), 653–82.
- Del Pinal, Guillermo (2015). Dual Content Semantics, Privative Adjectives, and Dynamic Compositionality. *Semantics & Pragmatics*, 8(7), 1–53.
- Del Pinal, Guillermo (2018). Meaning, Modulation, and Context: A Multidimensional Semantics for Truth-Conditional Pragmatics. *Linguistics and Philosophy*, 41, 165–207.
- Diaz-Legaspe, Justina (2019). What Is a Slur? *Philosophical Studies*, 42(2), 1–24.
- DiFranco, Ralph (2015). Do Racists Speak Truly? On the Truth-Conditional Content of Slurs. *Thought*, 4, 28–37.
- Dummett, Michael (1973). *Frege: Philosophy of Language*. Duckworth.

- Hay, Ryan J. (2013). Hybrid Expressivism and the Analogy between Pejoratives and Moral Language. *European Journal of Philosophy*, 21(3), 450–74.
- Hom, Christopher (2008). The Semantics of Racial Epithets. *The Journal of Philosophy*, 105(8), 416–40.
- Hom, Christopher and Robert May (2013). Moral and Semantic Innocence. *Analytic Philosophy*, 54(3), 293–313.
- Hom, Christopher and Robert May (2018). Pejoratives as Fiction. In David Sosa (Ed.), *Bad Words Philosophical Perspectives on Slurs* (108–31). Oxford University Press.
- Hornsby, Jennifer (2001). Meaning and Uselessness: How to Think About Derogatory Words. *Midwest Studies in Philosophy*, 25, 128–41.
- Huddleston, Tom Jr. (2016, May 2). Misogynists Go Crazy Over The New *Ghostbusters* Trailer. *Time*. <https://web.archive.org/web/20221113223814/https://time.com/4314767/ghostbusterstrailer-negative-comments/>
- Jeshion, Robin (2013a). Expressivism and the Offensiveness of Slurs. *Philosophical Perspectives*, 27(1), 231–59.
- Jeshion, Robin (2013b). Slurs and Stereotypes. *Analytic Philosophy*, 54(3), 314–29.
- Jeshion, Robin (2018). Slurs, Dehumanization, and the Expression of Contempt. In David Sosa (Ed.), *Bad Words Philosophical Perspectives on Slurs* (77–107). Oxford University Press.
- Kim, Eun Kyung (2014, March 15). Obama on ‘Mom Jeans’: ‘I’ve Been Unfairly Maligned.’ *Today*. <https://web.archive.org/web/20221208032701/https://www.today.com/style/obama-momjeans-ive-been-unfairly-maligned-2D79379010>
- Lasher, Megan (2016, June 22). *Ghostbusters* Director Says ‘Chick Flick’ Is a ‘Derogatory’ Term. *Time*. <https://web.archive.org/web/20221113223823/https://time.com/4378542/ghostbusterschick-flick-interview/>
- Luu, Chi (2019, June 12). The Dubious Art of the Dad Joke. *JSTOR Daily*. <https://web.archive.org/web/20221204022510/https://daily.jstor.org/the-dubious-art-of-the-dad-joke/>
- Manne, Kate (2017). *Down Girl: The Logic of Misogyny*. Oxford University Press.
- McCready E. (2010). Varieties of Conventional Implicature. *Semantics & Pragmatics*, 3(8), 1–57.
- McKinnon, Rachel (2018). The Epistemology of Propaganda. *Philosophy and Phenomenological Research*, 96(2), 479–89.
- Neufeld, Eleonore (2019). An Essentialist Theory of the Meaning of Slurs. *Philosophers’ Imprint*, 19(35), 1–29.
- Nunberg, Geoff (2018). The Social Life of Slurs. In David Sosa (Ed.), *Bad Words Philosophical Perspectives on Slurs* (237–95). Oxford University Press.
- Pustejovsky, James (1995). *The Generative Lexicon*. MIT Press.
- Richard, Mark (2008). *When Truth Gives Out*. Oxford University Press.
- Rosch, Eleanor (1978). Principles of Categorization. In Eleanor Rosch and Barbara B. Lloyd (Eds.), *Cognition and Categorization* (27–48). Lawrence Erlbaum.
- Rosch, Eleanor and Carolyn Mervis (1975). Family Resemblances: Studies in the Internal Structure of Categories. *Cognitive Psychology*, 7, 573–605.
- Saka, Paul (2007). *How to Think about Meaning*. Springer.
- Sperber, Dan and Deirdre Wilson (1995). *Relevance: Communication & Cognition* (2nd ed.). Blackwell.

- Tirrell, Lynne (1999). Derogatory Terms: Racism, Sexism and the Inferential Role Theory of Meaning. In Kelly Oliver and Christina Hendricks (Eds.), *Language and Liberation Feminism, Philosophy and Language* (41–79). SUNY Press.
- Whiting, Daniel (2013). It's Not What You Said, It's the Way You Said It: Slurs and Conventional Implicatures. *Analytic Philosophy*, 54(3), 364–77.
- Williamson, Timothy (2009). Reference, Inference, and the Semantics of Pejoratives. In Joseph Almog and Paolo Leonardi (Eds.), *The Philosophy of David Kaplan* (137–59). Oxford University Press.
- Wilson, Deirdre and Dan Sperber (2006). Relevance Theory. In Laurence R. Horn and Gregory Ward (Eds.), *The Handbook of Pragmatics* (607–32). Blackwell.
- Wilson, Deirdre and Robyn Carston (2007). A Unitary Approach to Lexical Pragmatics: Relevance, Inference and Ad Hoc Concepts. In Noel Burton-Roberts (Ed.), *Pragmatics* (230–59). Palgrave-Macmillan.
- Zeman, Dan (2021). A Rich-Lexicon Theory of Slurs and Their Uses. *Inquiry: An Interdisciplinary Journal of Philosophy*, 65(7), 942–66.