SYMPOSIUM

MORAL AND SEMANTIC INNOCENCE

CHRISTOPHER HOM
Texas Tech University

ROBERT MAY
University of California, Davis

1. Introduction

Moral Innocence is a term meant to describe that facts like the following obtain: that no Jews are kikes, that there are no kikes, but that there are Jews. It is the view, to be more prosaic, that the world we live in contains no such things as kikes, niggers, or chinks, but that it does contain Jews, African-Americans, and Chinese. These facts are the contents of the thoughts that no Jews are kikes, that there are no kikes, but that there are Jews; the thoughts that are expressed by the sentences:

(1) No Jews are kikes
(2) There are no kikes
(3) There are Jews

That sentences (1)–(3) are jointly true we refer to as *semantic innocence*, and it is from this perspective that we will approach moral innocence.

Moral innocence, on our view, is moral realism applied to pejoratives. The justification of this view, we will argue, stems from the simplicity of the explanation that semantic innocence affords for the epistemic query that is fundamental to any account of pejoratives:[1]

How can a competent, rational speaker of a language know the meaning of a pejorative without being committed to, or even complicit with, racist attitudes?

Our goal is to characterize the meanings of pejorative words, the truth-conditions of sentences that contain them, and the corresponding semantic

---

1. Although the term "racist" applies to race-based hateful attitudes, we will use it more generally to cover hateful attitudes towards religious, ethnic, national, and other such group-based identifications.

knowledge of competent speakers of a language. The foundation to our analysis rests upon a novel articulation of pejorative words themselves, and their connection to certain, *a priori*, moral facts.

Before embarking however, there are two matters on which we must insist if we are to stay on course. The first is that we must clearly distinguish *slurs*, as (parts of) token speech acts from *pejoratives* as linguistic expressions that may be deployed in those speech acts. Slurs are the focus of the act of slurring. Our view is exclusively about pejoratives as linguistic expressions, and their underlying linguistic properties. Making this distinction is crucial in understanding the project as a genuinely linguistic enterprise, and one that investigates the nature of linguistic meaning as opposed to the nature of language use. The second is that we are not concerned with the evolution of word-forms, with the criteria by which they come to be associated (or disassociated) with pejorative content. These criteria, it seems to us, are most fruitfully thought of as determined by social convention, and this is not a point of disputation in this paper.[2] All that we require synchronically is that there are pejorative terms, and our goal is to explore paradigm instances in order to characterize the essential linguistic properties of words of this sort.

2. Null Extensionality Thesis

We begin with an obvious observation, namely that pejorative terms are typically paired with non-pejorative terms with which they are conceptually linked—"Jew" and "kike" are paradigmatic in this regard—and an equally obvious question as to what is the nature of this relation? Fortunately, the answer to the question is straightforward, as we have a standard theory of the conceptual relations that hold of the sort of general terms instanced by "Jew" and "kike." This is the theory of generalized quantifiers.[3]

The theory of generalized quantifiers characterizes a class of relations R that apply to sets $X, Y \subseteq \mathscr{P}(D)$, where by taking X and Y to be extensions of concepts A and B, the clauses that govern the relations R determine conceptual relations. Of these relations, there is a core group that when taken together determine the "classical" relations between concepts, expressed by the following clauses:

$$\forall(X, Y) = \text{T } iff \ X \cap Y = X$$

$$\exists(X, Y) = \text{T } iff \ X \cap Y \neq \varnothing$$

$$\nexists(X, Y) = \text{T } iff \ X \cap Y = \varnothing$$

---

2. For example, we are not in disagreement with the sociolinguistic views presented in Brontsema (2004) and Lepore and Anderson (2013) that argue for the role of social convention in determining what words count as taboo.
3. For recent overviews, see Glanzberg (2006) and Antonelli and May (2012).

Concepts *A* and *B* are *universally* related if their extensions satisfy the first clause, *existentially* related if they satisfy the second clause, and *unrelated* if they satisfy the third. Applied to natural language, these clauses are standardly understood to specify truth-conditions, where these relations are expressed by the determiners "all," "some," and "no," respectively. Thus, "All men are mortal" is true if and only if the men are included with the mortals, "Some men are mortal" is true if and only if they intersect, and "No men are mortal" if and only if there is no intersection.

With just this much "off the shelf" semantics, we make the following observation: it directly follows that sentence (1) above—"No Jews are kikes"—is true, as there is no intersection of the class of Jews with the class of kikes, and accordingly both "Some Jews are kikes" and "All Jews are kikes" are false.[4] Or to put it a little differently, it follows that "Jew" and "kike" are conceptually unrelated. We elaborate.

The truth-conditions just described for sentence (1) entail that the extension of the word "kike" is not the same as the extension of the word "Jew"; more strongly, it entails that their extensions are disjoint. They are disjoint for the following reason: pejorative expressions like "kike," unlike "Jew," have null extensions. Hence not only is sentence (1) true (the null set is the intersection of any set with the null set), but so too are sentence (2)—"There are no kikes"— and sentence (3)—"There are Jews".[5] These cases illustrate the general thesis of *null extensionality*, i.e. that pejorative terms have empty extensions. Semantic innocence—the account of the truth-conditions of sentences containing pejorative terms—is a consequence of null extensionality.

Why do pejorative terms have null extensions? Why is nobody a kike? The reason is this: there are no morally evaluable traits (good or bad) that are heritable on the basis of race, gender, sexual orientation, and the like.[6] Accordingly there can be no terms that are satisfied in virtue of there being individuals having those traits. There are no kikes because there is no one who ought to be the object of negative moral evaluation just because they are Jewish. More generally, *no one* ought to be derogated for such reasons; *no one* is the target of pejoration. Nevertheless, it is part of the meaning of "kike" that Jews are the intended target of pejoration, and it describes this target in particularly negative and egregious ways. But pejorative terms so radically *mis*describe their

---

4. A common reaction to the negative existential statement is that it is really a case of *metalinguistic negation*; cf. Horn (1989, Ch. 6). While we suspect that metalinguistic negation is actually a heterogeneous phenomenon, Horn is clear that one diagnostic for metalinguistic negation is that it does not license negative polarity items. Negative polarity items, like "ever," are expressions that can only occur in downward entailing contexts, specifically in the scope of a negation. The following examples show that the proposal fails Horn's diagnostic: "Max is a kike. Don't speak that way. No Jews are *ever* kikes" versus "No Jews are *sometimes* kikes," which is ungrammatical in the context. Of course matters here are likely more complex, but this initial failure to meet Horn's diagnostic for metalinguistic negation casts doubt on the analysis.

5. This follows on the singular counterparts of the clauses in the text; i.e. $\nexists(X) = T$ *iff* $X = \emptyset$ and $\exists(X) = T$ *iff* $X \neq \emptyset$, which interpret "There are no *A*" and "There are *A*," respectively.

6. By "heritable" we do not mean in an evolutionary sense. Rather we mean inheritance in virtue of group membership, regardless of whether membership in that group is voluntary or not. Converts to Judaism are just as much targets of anti-semitism.

targets that they can have no extensions. Someone who believes otherwise, that there are kikes, someone who is in the grip of the pathology of racism and discriminatory beliefs, is either wrong about the world, or they do not know the meaning of the word "kike." That pejorative terms have null extensions is thus a morally significant fact that sharply distinguishes the meanings of pejoratives from those of their characteristic counterpart terms.

Accordingly, it is a consequence of null extensionality that paired terms like "kike" and "Jew" differ not only extensionally, but also intensionally; they have different meanings. Framing this in a Fregean perspective, since "kike" and "Jew" have different extensions, it follows that they must refer to different concepts; but this entails that they must express different senses, since sense determines reference.[7] As its reference, the sense of "kike" determines a concept with a null extension, while the sense of "Jew" determines a reference whose extension is non-null.[8] "Jews are kikes" is false,[9] and any proposition (thought) expressed by a sentence of the form $\ulcorner \Phi(kike) \urcorner$ is different from that expressed by a sentence of the form $\ulcorner \Phi(Jew) \urcorner$. Null extensionality thus implies that pejorative terms and their characteristic counterparts express different linguistic meanings, and as such they also reflect different and compatible competency requirements on the part of speakers. What one knows in virtue of knowing the meaning of "kike" is not the same as what one knows in virtue of knowing the meaning of "Jew."[10]

3. The Epistemic Implication

We take it as given that the linguistic meaning for ideological terms (inclusive of pejoratives) is a distinct (though closely related) entity from the ideology itself. Just as knowing the meaning of terms like "Nazi" or "slavery" doesn't make a speaker complicit with either associated ideology, knowing the meaning of a pejorative does not make one complicit with its associated discriminatory attitude. Why might this even be in doubt? One reason is that because the derogatory potential for such words is so strong, any association with the word conventionally signals an association with its corresponding discriminatory attitudes. Moreover, these attitudes are so negative that without an *outright repudiation* (i.e. the rejection of the term from one's lexicon), there is an available interpretation by hearers of the implicit adoption of those attitudes by the speaker. Since knowing the meaning requires having the term in one's lexicon, there is at least a suggestion of complicity and so knowledge of meaning *allows for the possibility* that its speaker is complicit with racism. Of course, for racists,

---

7. This manner of speaking makes clear that pejorative terms have references, *viz.* to functions that map all arguments to the False. It would thus be a mistake to say that pejorative terms would be non-referential if they have null extensions; cf. Williamson (2009: 149), who mistakenly assumes that in order for pejoratives to be referential they must have non-null extensions.

8. That the extension of "Jew" is non-null is contingent, but it is not contingent that the extension of "kike" is null. See discussion in section 4.

9. Read either as an identity or a predication. Note that "chinks are kikes" is true. It is true vacuously in the way that "dragons are unicorns" is true.

10. See Hom (2008: 430–432) for development of the intensional aspect.

296

this possibility is actualized, while for non-racists, it is a mere possibility. At the same time, it appears that racists and non-racists successfully understand each other with uses of such terms. Hence, the recommendation for *silentism* in a strong form—the complete rejection of racial pejoratives from the lexicon—is mostly driven by prudential reasons, and not by any deep theoretical view about their meaning.

This is not our view. Our view is that actual knowledge of the linguistic meaning of a pejorative term paired with *a priori* knowledge of moral facts entails knowledge of its null extensionality. From this it straightforwardly follows that one can know the meaning of a pejorative without being committed, even implicitly, to racist attitudes: if one knows certain *a priori* facts such as that being of a race or religion is not morally evaluable, then knowing the meaning of a pejorative is sufficient to know that the concept it expresses is not instantiated. Hence the fundamental epistemic issue posed at the outset is settled by adopting null extensionality.

Concomitantly, null extensionality is central in explaining the defective nature of racist claims as misrepresenting their targets, as well as the underlying pathology involved in making such claims.[11] Racists believe that pejorative terms have non-null extensions; they believe that there are kikes, chinks, and the like. They think that "No Jews are kikes" is false, because at least some Jews, perhaps even all, are kikes. But they are wrong, and they are wrong about what the word "kike" means. They may grasp the same underlying concept as a non-racist, yet nevertheless they are linguistically incompetent as that knowledge does not issue the right truth-conditions for sentences containing pejoratives. In contrast, non-racists do properly know the meanings of pejorative terms, and assign the proper truth-conditions. They know that there are no such things as kikes, niggers, and chinks; that there are no people who are the proper objects of pejoration. Because non-racists know the meanings of pejorative terms and properly grasp the moral concepts that they express, they will show caution in using them. In contrast, racists will throw caution to the wind, since they think there are people who are properly described as the objects of morally negative evaluation, just because of being a member of a certain group (i.e. for being a Jew).[12]

We are inclined at this point to take matters as settled, and take moral innocence as vindicated, rejecting any alternative account that rejects null extensionality. The reason is simple. If the linguistic meanings of pejoratives entail that they have non-null extensions, then this implies that there are kikes, niggers, and chinks, and accordingly that racists' beliefs are true and justified. To our mind, any theory with this implication is unacceptable, and we hardly

11. In section 8, we return to this issue to show how this pathology is difficult to account for without null extensionality. The problem with views that deny null extensionality is that they mistakenly attempt to locate the pathology of racist pejoratives as non-central to their truth-conditional, semantic contents.

12. This morally negative evaluation encompasses a wide range of negative potential attitudes from hate and derision on one extreme to mere dislike or resentment on the other. The "skinhead anti-semite" and the "country club anti-semite" express the same semantic content with their uses of "kike," though their surrounding dispositions and psychologies are likely to vary tremendously.

think anyone would disagree. Nevertheless, it might be thought that we are being too hasty in our rejection; assuredly, we are warranted in issuing a cautionary note, but adjudicating whether the objection is fatal requires a more careful examination both of null extensionality and of alternatives. We agree, and so we continue the discussion, to both articulate our view of the meaning of pejorative terms, as well as to clarify our objections to other accounts.

4. The Lexical Analysis of Pejoratives

Pejorative terms are underlyingly complex lexical items. They are composed of a lexical marker of pejoration—PEJ—which combines with a relevant characteristic counterpart term, yielding a fully-formed, overtly pejorative expression. For example, "kike" is lexically analyzed as PEJ(*Jew*), where PEJ is a lexically covert marker of pejoration that functionally combines with the term "Jew." More generally, PEJ(ξ) functionally combines with any characteristic counterpart term, t, typically (but not exclusively) designating race, gender, religion, class, and so forth, to form a pejorative, PEJ(t).[13]

While PEJ(ξ) is unambiguous at the lexical level, its phonological realization allows for significant variation. This explains why pejorative language can be so exceptionally difficult to interpret on the part of its audience. Consider some of the various dimensions of phonological realization for pejoratives. The base cases are the clear, overtly recognized pejoratives (e.g. "kike," "chink," "nigger"). In addition, there are assorted pejorative modifiers (e.g. "dirty," "rotten," "one of *those*") that take characteristic counterpart terms as inputs (e.g. "Jew," "Chinese," "African American") to generate pejorative terms (e.g. "dirty Jew," "rotten Chinese," "one of *those* African Americans"[14]). Finally, there are extralinguistic markers like tone (e.g. sneering) and gesture (e.g. looks of disgust) that can accompany utterances of characteristic counterparts to generate pejoratives. We take these all to be conventionalized means of expressing negative bias, and we postulate that they are lexically marked at a more abstract level by PEJ(ξ).

The semantic interpretation of PEJ(ξ) is neatly expressible in Fregean terms. The sense expressed by PEJ denotes a second-level function that combines with a first-level concept, (e.g. of race, gender, religion, or class) to form a complex first-level pejorative concept. The complex pejorative concept takes objects as inputs, and has falsity as its output. That complex pejorative concepts are constant functions mapping individuals to falsity is the realization of the null extension thesis. Intensionally, we can think of PEJ as a concept abstracted from the moral truth that nobody ought to be the target of negative moral evaluation because of being Jewish; that is, the concept *x ought to be the target of negative moral evaluation because of being* ξ. This concept can be abstracted from any

---

13. Note that we mean "characteristic" here in the sense of characteristic function, so that semantically the argument term of PEJ serves solely to pick out a class of entities.
14. Note that this analysis of the demonstrative pejorative modifier helps to account for those racists who hold that only some proper subset of the target class are despicable for racist reasons.

truth of this form, that is, for any substitution for "Jew" *salva veritate*; accordingly, each of the distinct instances of PEJ(ξ) will correspond to a distinct pejorative term, and each will have a null extension.[15] That pejorative terms are necessarily unsatisfiable can be illustrated by comparing PEJ to other modifier terms that work in a similar fashion to drain the extension of the terms on which they operate. Examples include such modifiers as "fictional" and "magical." When combined with ordinary referential terms like "car" or "horse," the modified expression no longer retains an extension. There are cars, but no *fictional* cars. There are horses, but no *magical* horses, i.e. there are no unicorns. In similar fashion, there are Jews, but there are no PEJ(*Jew*), i.e. there are no *kikes*.

"Magical" provides a particularly useful comparison to PEJ because they each bring to bear their own false conceptual point of view. In the case of "magical," the term is ontologically and supernaturally loaded with false assumptions about the causal structure of the world. Nothing is magical, and so "magical," as a modifier term, drains extensionality. Similarly PEJ is ideologically loaded with false assumptions about the social and ethical structure of the world. No one ought to be the target of derogation because of their particular group membership, and so PEJ, as a modifier term, also drains extensionality. Both modifiers draw their capacity from false assumptions that conflict with reality.

There is also an important distinction to be made between "fictional" and "magical" that highlights another important feature of PEJ. Whereas "fictional" wears its falsity on its sleeve, "magical" does not. Normal fictional uses of language presuppose that the audience shares a mutual recognition with the author that the language is not truth-seeking; i.e. that fictional terms are not about real objects. This feature need not hold for "magical"; e.g. when people actually believed in magic. We take this to be the threshold distinction between fiction and propaganda. Where fiction shares this feature of mutual recognition, propaganda is a social or political kind of story-making that authors engage in to propagate false beliefs among their audience. Users of these empty ideological terms intend that their audience *misrecognize* the false nature of this language, often with the further intention to manipulate their behavior. The goal is to inflame, not inform. In this way, proponents of propaganda speak with bad faith. Whether they are committed to the truth of their claims is irrelevant for them and their primary goal.

The disingenuous quality of pejoratives that stems from their origin as terms of propaganda leads to the further result that null extensionality follows *necessarily*. In the postscript, Kripke (1980) argues that "unicorn" has a necessarily

---

15. Terms for race, gender, religion, and class are clear cases of substitutions *salva veritate*, but bear in mind that we are not assuming that this class is universal. We allow that there may be substitution instances that do not preserve truth, e.g. that nobody ought to be the target of negative moral evaluation because of being a Nazi or a child molester, although we recognize that someone who holds that we should abjure all hateful attitudes would disagree. But if there are such instances, then PEJ is not abstractable from them, and hateful words directed at these groups would not be pejoratives. Accordingly, they may be socially sanctioned in ways that pejoratives are not.

empty extension because there were no actual unicorns available at the original dubbing. So nothing actually shares in the causal relation to the name. For him, the myth is indeterminate to specify whether any possible animal that shares particular descriptive properties associated with unicorns are genuinely unicorns. For these reasons, he concludes that there *could not* be unicorns. The same line of reasoning follows for "kike" as an instance of PEJ(*Jew*). Given that racist ideologies are radically false, there are no referential entities that fit the originating speaker intentions for these terms. Hence the ideology is indeterminate to specify whether any possible object that shares particular descriptive properties associated with the pejorative genuinely falls in its extension. For these reasons, nothing can fall inside the extension of a pejorative term. Recall that proponents of propaganda act on bad communicative faith, so the stronger claim for the necessity of null extensionality is completely beside the point given their overall political aims.

Although the null-extensionality of "kike" and "unicorn" are both necessary, they nevertheless differ in that this is knowable *a priori* only for the former, as we saw in the previous section. The reason for this is the tight connection between the meanings of pejoratives (e.g. that these people ought to be treated negatively on the basis of their race) and the moral structure of the world (e.g. that no one ought to be treated negatively on the basis of their race). That terms like "unicorn" or "magic" have null extensions is, in contrast, *a posteriori*. To know whether there are magic horses requires empirical investigation, and this is because the meaning of "magic" does not share in a sufficiently tight connection to the world that it describes. In this regard, words like these differ from pejorative words.

5. The Compositionality of PEJ

Intrinsic to our abstract lexical characterization of PEJ is that it may be phonologically realized in a variety of ways. Some of these may directly reflect the underlying lexical composition, for example, "dirty Jew" as a realization of PEJ(*Jew*), but not necessarily; "kike" is lexically non-distinct from "dirty Jew" or from "*Jew*" uttered with a sneering tone. All of these are instances of the same underlying lexical item.[16]

The compositional aspect of PEJ(ξ) explains some of the seemingly ambiguous interactions that surround pejoratives and amplifying modifier terms. Consider the following sentences and their readings:

(4) John is a terrific Jew.
    (a) John is very observant.
    (b) John is a paradigm example of anti-semitic stereotypes.
    (c) John is BAD for being Jewish.

16. Because of the functional characterization of PEJ(ξ), it can iterate. We can have PEJ (PEJ (t)), PEJ (PEJ (PEJ (t))), and so on. This reflects in expressions like "dirty Jew," "dirty rotten Jew," "dirty rotten kike," etc. Unsurprisingly, all these display a certain redundancy of meaning; it would be like saying "my best best friend."

300

(5) John is a horrible Jew
   (a) John is not very observant.
   (b) John is a bad example of anti-semitic stereotypes.
   (c) John is BAD for being Jewish.

The *a*-readings require a non-pejorative, neutral reading of "Jew," and in this regard contrast with the *b*- and *c*-readings, which require a pejorative reading of "Jew." The *b*- and *c*-readings, in turn, are distinguished; only the latter is equivalent to "John is a kike."

These observations follow directly, depending upon whether "terrific" and "horrible" occur as amplifying adjectives, or as phonological manifestations of PEJ. We can then distinguish the cases as follows:

(4a') $\lambda x[\text{Terrific (Jew)}](x)(\text{John})$
(4b') $\lambda x[\text{Terrific (PEJ(Jew))}](x)(\text{John})$
(4c') $\lambda x[\text{PEJ(Jew)}](x)(\text{John})$
(5a') $\lambda x[\text{Horrible (Jew)}](x)(\text{John})$
(5b') $\lambda x[\text{Horrible (PEJ(Jew))}](x)(\text{John})$
(5c') $\lambda x[\text{PEJ(Jew)}](x)(\text{John})$

The difference between the *a*- and *b*-cases is that only in the latter does the amplifying adjective apply to a pejorative term. The *c*-case is like the *b*-case, as it too contains a pejorative term, but it does not contain an amplifying modifier. Rather, this is the case in which "terrific" and "horrible" function as overt realization of PEJ. The analysis, we note, generalizes to a broader class of modifiers, "amazing," "awful," "incredible," "wonderful," "extreme," and so forth.

That pejoratives are underlyingly compositional also explains the apparent lexical ambiguity surrounding cases of appropriation. Appropriation is the complex phenomenon whereby targeted groups seek to take control and modify pejorative words for their own political and social purposes.[17] On our view, appropriation is elegantly modeled as the lexical removal of PEJ at this more abstract level of lexical analysis. So even though the surface realization may remain the same, the underlying expression has shifted in the appropriated case to cancel pejoration. Appropriation is basically the process whereby targeted groups seek to modify the syntactic identity of pejorative terms. Because there is no overt marker for this shift, appropriated uses of a token expression (e.g. "nigger" or "nigga") are difficult to distinguish from pejorative uses; i.e. their lexical identity is determined at a more abstract level, and hence the covert nature to these terms. The covertness of PEJ(ξ) thus offers an intuitive account for the difficult aspects of recognizing the appropriation of pejoratives.

17. For a detailed analysis, see Brontsema (2004).

6. The Aboutness Objection

It might be thought that there is an objection to our point of view; i.e. if pejoratives have null extensions, then sentences that contain them are not about anything. Yet surely racist users of such words manage to say (and are understood to say) derogatory things about something, namely members of their target class—after all, isn't someone who utters "Kikes are usurious" saying something about Jews? Call this *The Aboutness Objection* to the thesis of null extensionality.

We contend that the objection comes about because of a fundamental misunderstanding regarding the linguistic nature of pejoratives. We can get a sense of this misunderstanding from the following remark of Frege's, where he distinguishes *concepts* from their *characteristic* marks:

> If one says: "A square is a rectangle in which adjacent sides are equal", then one defines the concept *square* by stating what properties something must have in order to fall under it. I call these properties characteristic marks of the concept. . . . Whether there are such objects [falling under the concept] is not immediately known on the basis of the definition . . . nor does the definition guarantee that the concept is instantiated. (Frege 1893: XIV)

Pace Frege, it is not sufficient for a concept to be instantiated that its characteristic marks are instantiated; even if the characteristic marks of a concept are instantiated, it does not follow that the concept is. This is the case with pejoratives. *Jew* as it occurs in PEJ(*Jew*) is a characteristic mark of the pejorative concept; although "Jew" has an extension, "kike" does not. This circumstance, however, does not impede a characteristic mark from playing its conceptual role of setting the essential appropriateness conditions for the application of the concept. Thus, for pejoratives, the characteristic marks fix the target of pejoration. It makes PEJ(*Jew*) be *about* Jews, and not about some other group. For this reason, there is an accentuated cognitive dissonance when incorrectly applying pejorative terms outside of their intended target class. To call a gentile a "kike" is not only false, but also conceptually incorrect, as gentiles fall outside of the characteristic mark for the concept PEJ(*Jew*).

Statements that contain pejorative terms are about their intended targets because those targets satisfy the characteristic marks of pejorative concepts, not because they fall under those concepts. Hence, the aboutness objection may be set aside.

7. Moral Corruption

With our positive view in place, let's consider the alternative to moral innocence. Without bias, call this view *Moral Corruption*. Such a view denies null extensionality, holding instead that the extensions of pejoratives and their neutral counterparts overlap. According to corruption, there are kikes, because there are Jews. What are the consequences to corruption and how are we to analyze the corresponding beliefs of racists and non-racists under such a view?

302

One option for moral corruption is to hold that racists have true beliefs about the existential facts (e.g. that there are kikes), and that non-racists have false beliefs (e.g. that there are no kikes), end of story. We take this view as perniciously a non-starter. This would make racism true, not false, and would lead to the absurd result that racists have deeper insight into the world than non-racists.

Another option for moral corruption is to hold that both racists and non-racists hold identically true beliefs about the facts. The fact that there are Jews *just is* the fact that there are kikes. These facts are identical because the semantic content of "Jew" is identical to the semantic content of "kike," and hence, it will be necessary that all Jews are kikes. In other words, the underlying racism that the non-racist objects to, (i.e. the negative moral content associated with pejoratives), can not lie with the truth-conditional contribution of the pejorative term, but rather with something else. This is the only feasible option for corruption, which holds that pejoratives and their characteristic counterparts are truth-conditionally synonymous.

8. Difficulties Relocating Pejorative Significance

Let us call the plausible version of moral corruption just described the *identity thesis*. On this account, the truth-conditional contents of pejoratives and their "neutral" counterparts are identical; that is, the Jews = the kikes. Since pejoratives and their neutral counterparts differ in their pejorative significance, identity theorists must postulate some further kind of non-truth-conditional content that distinguishes pejorative terms from their counterpart terms. Call this *expressive content*. Expressive content is supposedly non-truth-conditional, and functions as the primary explanation for what is bad about pejoratives. We won't go into details here, as there are a family of articulable views.[18] The essential character of the genus is that expressive content non-truth-conditionally captures the pejorative aspect of pejorative terms. Call the resulting conjunction of these two hypotheses, *identity-expressivism*. We contend that identity-expressivism faces four major challenges that makes Moral Corruption an unappealing alternative to Moral Innocence.

8.1 The Problem of Contrasting Pairs

First, the off-loading of pejorative significance from truth-conditions is *prima facie* problematic because of how sentence pairs like the following behave under substitution:[19]

---

18. For example, see Kaplan (1999), Copp (2001), Hornsby (2001), Potts (2005), and Williamson (2009).
19. Examples from Hom (2008).

Institutions that treat Jews as <u>kikes</u> are anti-semitic. (True)
Institutions that treat Jews as <u>Jews</u> are anti-semitic. (False)
Anyone who thinks that Jews are <u>kikes</u> is anti-semitic. (True)
Anyone who thinks that Jews are <u>Jews</u> is anti-semitic. (False)
I'm not a <u>kike</u>, but I am a Jew. (True)
I'm not a <u>Jew</u>, but I am a Jew. (False)
Max doubts that Jews are <u>kikes</u>. (True)
Max doubts that Jews are <u>Jews</u>. (False)
Thinking that Jews are <u>Jews</u> is to be radically wrong about the world. (False)
Thinking that Jews are <u>kikes</u> is to be radically wrong about the world. (True)
<u>Kikes</u> are supposed to be Jews that are bad. (True)
<u>Jews</u> are supposed to be Jews that are bad. (False)

Because each sentence of the pair differs only in the inter-change of the pejorative and its non-pejorative counterpart, their resulting difference in truth-value is evidence that pejoratives make truth-conditional contributions to sentences in which they occur. Critically, those truth-conditional contributions must be derogatory for the correct result in the overall truth-value of the sentence. This indicates that postulating expressive content is superfluous. The difficulty is reinforced by consideration of pairs of questions such as the following that appear to express a different range of acceptable answers:

Am I anti-semitic if I think that there are <u>kikes</u>? (Yes)
Am I anti-semitic if I think that there are <u>Jews</u>? (No)
Why do anti-semites think that Jews are <u>kikes</u>? (ignorance, hatred, etc.)
Why do anti-semites think that Jews are <u>Jews</u>? (analyticity, rationality, etc.)

Again, the only relevant difference to each member of the pair is the switch of the pejorative and its counterpart term. The resulting shift in the conditions for acceptable answers again suggests that the pejorative makes a derogatory, truth-conditional contribution to the question asked, and hence making expressive content superfluous.

8.2 Modal-Conceivability

The second problem is that there is a version of the modal-conceivability argument against identity-expressivism that runs as follows:[20]

1. It's conceivable for there to be Jews without kikes.
2. Whatever is conceivable is possible.
3. Therefore it's possible for there to be Jews without kikes.

The conclusion to the modal-conceivability argument is inconsistent with identity-expressivism which is committed to the identity of the extensions of "Jew" and "kike." To deny the conclusion, the identity-expressivist must deny

20. The argument is modeled after Chalmers (1996).

304

premise 1. Why think that premise 1 is true? While we think that the actual world is sufficient to illustrate the truth of premise 1, one can also consider a world that is mostly like ours, except that it is *morally perfect*, and so devoid of racism of any kind. In such a world, it is clearly conceivable for there to be Jews without kikes. What is the identity-expressivist's response? One likely response is to offer a reinterpretation of premise 1 and the conclusion such that identity-expressivism is preserved. Notice that on their view, premise 1 literally should read as:

1'. It's conceivable for there to be Jews without Jews.

But on this reading, premise 1 is inconsistent, and hence the conclusion necessarily follows. So the identity-expressivist must offer some other reinterpretation of the argument, such as:

1". It's conceivable for there to be Jews without Jews towards whom hatred is conventionally directed.
2. Whatever is conceivable is possible.
3". Therefore it's possible for there to be Jews without Jews towards whom hatred is conventionally directed.

On this recasting of the modal-conceivability argument, the morally perfect world that might motivate someone to think that premise 1 is true is reinterpreted as support for premise 1" instead; i.e. the morally perfect world is where there are Jews but where the conventional expression of anti-semitism does not exist. This result is consistent with the moral perfection of the motivating example. More importantly, the reinterpretation of the conclusion is now consistent with identity-expressivism.

The problem with this reinterpretation is that it must explicitly codify the negative, expressive content such that it can do inferential work in the argument. Notice, too, that there *must* be a reinterpretation on the part of the identity-expressivist because of the inconsistency of 1'. But to codify the expressive content to do inferential work is simply to admit that expressive content is part of the truth-conditional contribution of the pejorative term. So like the examples in the previous section, the modal-conceivability argument provides further evidence that expressive content is superfluous.

8.3 A Frege–Geach Problem

The third problem facing identity expressivism is a version of the Frege–Geach problem, namely that in embedded contexts, pejoratives can successfully interact with truth-conditional and intensional operators, indicating that pejoratives make truth-conditional contributions to sentences in which they occur. The problem is crystalized under the plausible assumption that non-truth-conditional expressive content of a sort similar to that which supposedly explains the derogatory nature of racist pejoratives, also explains the derogatory nature of non-racist pejoratives. The non-racist pejoratives are useful for

305

this argument because they are free to embed in a wide variety of scoping contexts. This greater syntactic flexibility of non-racist pejoratives affords them a clearer, albeit less direct, test for variable embedding. In the following examples, the derogatory element does not scope away, as predicted if the derogatory element were non-truth-conditional, and hence not available for truth-conditonal and intensional operators:[21]

Negation
John didn't fuck the managing partner's daughter. (He hasn't even met her.)
John is not a bastard. (He's extremely nice.)

Conditionalization
If John fucks up another case, then he will be fired. (But I don't think he will because he's working much harder now.)
If John is a bastard, then the firm will not hire him. (But John is an extremely nice fellow.)

Indirect reports
My father told me that I could not marry that damn Brad.[22]
I'm not prejudiced against Caucasians. But John, who is, [thinks/claims] that you are the worst honky he knows.[23]

Event quantification
Every time someone fucks up a case, the senior partner blames John. (But as long as I'm not blamed for it, I don't care.)

Tense
John was a fucker in law school, but he has improved since then. (I like him quite a lot now.)

The successful interaction between pejorative and operator is further evidence that pejoratives make truth-conditional contributions to sentences in which they occur. The felicity of the parenthesized follow-up remarks also highlights that their truth-conditional contributions are indeed pejorative, and this is further reason to take expressive content as superfluous.

8.4 A Frege Puzzle

The fourth problem for identity expressivism is a version of Frege's Puzzle. Frege's Puzzle is a puzzle about how sentences of the form "$a = a$" are trivial and knowable *a priori* while sentences of the form "$a = b$" are non-trivial and knowable *a posteriori*, when a and b have identical referents. Identity expressivists face an instance of Frege's Puzzle with pairs like: "Jews are Jews" and "Jews are kikes." Because identity expressivism posits that "Jew" and "kike" have identical semantic values, they cannot appeal to the traditional

21. Uncited examples from Hom (2012).
22. Kratzer (1999: 6).
23. Schlenker (2003: 98).

solution offered by Frege (i.e. a difference in sense) to account for the difference in their cognitive significance.[24]

Identity expressivists, however, are not without recourse. Suppose that belief is a three-place relation (BEL) between agent, proposition, and mode of presentation, where a mode of presentation is a way of thinking of an object.[25] Under the analysis of belief as a three-place relation:

A believes that *p iff* ($\exists m$)BEL (*A, p, m*).

When an agent *A* thinks of Jews in an anti-semitic way, she employs one mode—call it m^K. When *A* thinks of Jews in a non-anti-semitic way, she employs another mode—call it m^J. The difference in cognitive significance between the identity pairs can now be naturally explained. On the one hand, "Jews are Jews" is trivial and knowable *a priori* because it expresses the proposition that Jews are Jews, and is believed under a complex mode of presentation that takes the form: $<m^J, m^J, m^=>$. Because the agent conceives of the inputs to identity under the same mode of presentation, m^J, the agent *automatically* knows that the objects must be the same, and hence that the truth of the identity claim follows *automatically*. On the other hand, "Jews are kikes" is non-trivial and knowable *a posteriori* because while it expresses the same proposition that Jews are Jews, it can also be believed under a distinct complex mode of presentation that takes the form: $<m^J, m^K, m^=>$. Because the agent conceives of the inputs to identity under distinct modes of presentation, $m^J \neq m^K$, it is conceivable that those modes represent distinct objects, and hence it is an open question with regard to the language that the truth of the identity claim follows. While the proposition may be true, its truth does not follow *automatically* for the agent as it did in the previous case. The difference in modes explains the potential epistemic difference for the agent, and hence the difference in overall cognitive significance.

At this point, the identity-expressivist takes the puzzle as solved, but we purport that problems arise. Consider a potential agent *B* who is herself an identity-expressivist, rational, competent, and a non-racist. As a competent speaker of English and an identity theorist, our agent believes that "Jew" and "kike" make the same truth-conditional contribution, and hence are terms that apply to the exact same class of individuals. As a rational agent, *B* reasons according to Leibniz's Law (the indiscernibility of identicals), so any property held by Jews is held by kikes. As a non-racist, *B* refrains from thinking of Jews in anti-semitic ways. But then how can *B* have distinct modes of presentation m^J and m^K? For if the modes are distinct, they must be different ways of thinking of Jews; that is, they must differ on at least one anti-semitic property ascribed to Jews, as m^K is an anti-semitic way of thinking of Jews. But because *B* is committed to both Leibniz's Law and the identity thesis, she believes that whatever holds of Jews must hold of kikes, and vice versa. In other words, to view kikes negatively *is just* to view Jews negatively. So however the expressive

24. The remarks in this section are drawn from the more extended discussion in Hom and May (forthcoming).
25. For example, see Salmon (1986), Schiffer (1987), and Soames (1987).

content determines m^K, it doesn't seem like m^K is a concept that B can utilize. To grasp the mode of presentation for "kike" is to conceive of Jews *as the anti-semite does*. It is to have Jews presented to the mind as they are presented to the anti-semite. Thus, to accept this particular solution to Frege's Puzzle would require non-racists to think of people in the way that racists think of them. This is exactly Mark Richard's point when he says:

> The attitude—the contempt—*is part of what one thinks*. Furthermore, to have an attitude of contempt towards someone because of their race or ethnicity is, *inter alia*, to represent one's target in a certain way: as contemptible because of his race or ethnicity. Such a representation is incorrect: no one is contemptible for such a reason. (Richard 2008: 14)

We agree with Richard's assessment here.[26] The point is that a conventionalized, anti-semitic mode of presentation associated with "kike" leads to the incorrect result that competent (non-racist) speakers of pejoratives have pejorative, anti-semitic ways of thinking. What is required here prudentially is something like a cognitive version of silentism where anti-semitic ways of thinking are banned from one's conceptual architecture. Not only does this saddle identity-expressivism with the unenviable position of having to recommend "thought-policing," but it seems to leave non-racists without the conceptual resources to fully understand pejoratives. Thus, it does not appear that the identity-expressivist herself has access to distinct modes of presentation as a solution to Frege's Puzzle without being either irrational or schizophrenic.[27]

We conclude this section by pointing out that moral innocence faces no such Frege Puzzle. Since the terms "Jew" and "kike" make distinct truth-conditional contributions (i.e. "Jew" and "kike" express different senses), the identity claims in the anti-Semitic Frege Puzzle express different propositions. "Jews are Jews" is trivially true, but "Jews are kikes" is simply false. Hence the difference in their cognitive significance. End of story.

9. Concluding Remark

The world is morally innocent. That people ought to be treated in a racist way is not part of the fabric of reality. Of course people are, but those who do treat

26. We differ from Richard, however, with respect to its significance. For us, the unacceptability of racist claims is intimately related to the null extension thesis itself; pejoratives express empty concepts. Richard, on the other hand, holds that pejoratives are completely non-referential; pejoratives fail to express any concept whatsoever. He advocates a version of an expressivist theory where the misrepresentation of a pejorative "deprives what is said of truth" (Richard 2008: 26). Even stronger for Richard, it deprives what is said completely of truth-value. For Richard, the overall analysis is part of his broader argument against the traditional notion of truth that we take no stand on here. It's important to be clear about how our view is distinct from Richard's view. On Richard's view, pejoratives have sense, but lack reference (i.e. the compositional determinants for truth-value).

27. In Hom and May (forthcoming), we observe that neither taking modes of presentation as ways of *feeling* nor as ways of *talking* helps the matter.

people this way are in the grip of a pathologically false ideology, and harbor perniciously false beliefs about the world. Our language is semantically innocent; it mirrors a morally innocent world. The meanings of pejorative words and the sentences that contain them reflect that there are no kikes or chinks or niggers. These do not exist. We as rational agents know that the world is innocent, and so know that the meanings of pejorative words are loaded with ideologies that do not obtain. Innocence is not the absence of knowledge; it is the absence of sin.[28]

Appendix

The Persistence of Offensiveness

The central worry articulated in this volume for the semantic view presented in Hom [2008], and by extension for the view we present here, is the persistence of offensiveness through various constructions, in particular, negation, as racists at least seem to negate pejoratives without canceling their offensiveness. The argument purports to show that because offensiveness persists under negation, the semantic content of the pejorative is therefore *not* its source. For example, consider the following sentences:

(1) Yao is a chink.
(2) Yao is not a chink. He's a jap.
(3) There will never be a chink PM.

Jeshion notes:

> [1]'s offensiveness, claims Hom, resides in its asserting racist beliefs and agendas of Yao. The denial of these beliefs and agendas cannot itself be the source of [2]'s offensiveness,

and Whiting maintains that:[29]

> . . . except perhaps in certain pedagogical contexts, it appears typically to be no less derogatory to make *negative* claims using slurs than it is to make *positive* claims using them, while [Hom] suggests otherwise.

29. Whiting's sentence (3) is somewhat clearer, as Jeshion's sentence (2) is clouded by the speaker's derogation of Yao for being Japanese and the likelihood that if someone is racist toward one major group of Asians, they are racist toward all Asians. For now, set this worry aside.

The problem, so the argument goes, is that the effect of negating pejoratives would be to cancel their offensiveness, and hence turn them from being offensive to being inoffensive. But if the semantic content of pejoratives is *the* source of offensiveness, then it seems that there is no account of the force of pejoratives as used in (2) and (3).

This is plainly a bad argument. To see why, consider what we say about (1), (2), and (3). Semantically, it is this: (1) is false, the first clause of (2) is true, while the second is false, and that (3) is true. This follows from 'chink' having a null extension. Negation in these sentences functions just as it does with the pair "Yao is Chinese" and "Yao is not Chinese"—the former is true, and the latter is false, but in both the meaning of the term "Chinese" is the same. The only issue is whether Yao satisfies this predicate. Similarly, for (1) through (3), the meaning of "chink" is invariant across the examples; negation does *not* polarize the predication, turning its content from offensive to inoffensive. "Chink" has the very same meaning in all of the examples.

Accordingly, if offensiveness is part of the semantic content of pejorative terms, and semantic content is preserved under negation, then obviously offensiveness will be preserved under negation. But to be clear: we do not accept the antecedent. What we do accept is that part of the semantic content of pejoratives is negative moral evaluation, and that *derogation* is the application of that content to an individual. To say that Yao is a chink is to affirm that Yao ought to be subject to such derogation, and to say that Yao is not a chink is to deny that Yao ought to be subject to such derogation. Again, either way the meaning of the pejorative is unchanged.

In distinction from derogation, offensiveness, both giving and taking, is a psychological phenomenon, and can have many sources which one may seek to understand.[30] Offensiveness can be linguistically triggered, because when speakers use predicates, they typically conversationally implicate their commitment to the non-null extensionality of the predicate. Thus, unless a speaker uses great care in their speech, they are liable to make utterances like (2) or (3) that trigger the offensive conversational implicitum that there are Chinese people deserving of derogation, (even if Yao is not one of them).[31] Non-racist speakers take this moral to heart. Racist speakers, on the other hand, may have little or no qualms in this regard, since the implicature generated by their utterances would be consistent with their (false) beliefs. Hence their use of pejoratives is different, as they do not recognize the difference between the concept of being Chinese and the concept of being a chink. Racists are mistaken in how they attempt to deploy pejoratives, and this is a paradigm reflection of their flawed and odious conception of the world.

30. As Whiting explicitly notes, the conceptual distinction between derogation and offensiveness has been previously articulated in Hom (2008, fn. 25; 2010, fn. 5; and 2012, §4 and §5). See Camp [2013] for interesting suggestions as to cognitive causes and effects of offensiveness. Alternatively, one might hold that the causes are psychologically inscrutable, as do Anderson and Lepore [2013], and that all we can say is that they are socially taboo.
31. Note that "Yao is not a chink" may be uttered without generating offense, for example, when a non-racist utters "Yao is not a chink, because there isn't any such thing" in a pedagogical manner. But again, given the invitation to the implicature, and the consequences of it, a non-racist speaker must show deliberate care in selecting a context of utterance.

The offensiveness of the conversational implicatures generated from negated pejorative utterances is unsurprising upon careful attention to the interaction between the semantic analysis of Hom [2008], negation, and DeMorgan's law. To abstract away from the details, let the idealized semantic analysis of (1) be: "Fy ∧ Gy ∧ Hy," where 'F,' 'G,' and 'H' represent stand-in predicates for a complex array of externally determined prescriptions, stereotypes, and causal relations among them—collectively the characteristic marks of "chink"—and where 'y' is a name for Yao. The negation of (1) will be: "¬(Fy ∧ Gy ∧ Hy)," whose DeMorgan's equivalent is: "¬Fy ∨ ¬Gy ∨ ¬Hy." Crucially, notice that the negation of (1) is consistent with a wide range of distribution of information, from the negation of all of the conjuncts to the negation of simply one of the conjuncts.[32] Thus, a racist who utters (2) implies that one of the conjuncts fails to hold—namely the characteristic mark of being Chinese. But *not* the rest. Because they are racists, they are committed to negative, normative relations between discriminatory practices and the racial stereotypes that they falsely believe *actually hold* for the targeted group in question—even if the salient individual (e.g. Yao) isn't in that group. Unsurprisingly, such a thick, racist implicature is difficult, though not impossible, to cancel.

The Identity Thesis

As theories of *slurring* (or language *use*), we can agree with much of what is said by the other contributors to this volume. The social and psychological dynamics of the uses of pejoratives are complex, and will likely require a variety of tools for their complete understanding. But as the primary explanation for the linguistic meanings of these terms, these views ultimately fail because they take the identity thesis for granted. Are the Jews and the kikes one and the same? We argue that they are not, but in the present context this is the minority view; rather the consensus view is that they are to be identified. On this view, the extensions of 'Jew' and 'kike' are the same, and "Jews = kikes" is true. As we note, however, this identity thesis can be maintained only if pejorative terms are bereft of moral content, at least as far as their contribution to the truth-conditions of sentences in which they occur is concerned. The locus of the moral content is thus to be shifted to some non-truth-conditionally implicated notion of meaning (Camp, Whiting, Jeshion) or shifted out of the semantics altogether (Anderson & Lepore). If moral content were not displaced, it would be hard to see how to avoid complicity with racist attitudes.

In section 8 of our paper, we present a series of arguments against the identity theory. It fails, we note, to account for differences in truth-conditions of sentences that minimally contrast in the occurrence of pejorative terms (§8.1), for the modal conceivability of their being Jews without their being kikes (§8.2), for Frege–Geach problems (§8.3), and for Frege Puzzles of identity and propositional attitudes (§8.4). Taken together this makes in our view for a

---

32. Thanks to Justin Morton for making this suggestion.

compelling case. Nevertheless, given that philosophers never lack for cleverness, we do not doubt that there may be ways, to greater or lesser extent, of accounting for these problematic cases. But that is quite the point; these accounts all amount to *explaining away* the difficulties, but this is hardly to explain, especially given the alternative. Thus, from the perspective of Moral Innocence, none of these cases pose any problem at all, and rather follow directly from the theory.

It is perhaps worth emphasizing that we find the keenest difficulty for the identity thesis with the last of the cases mentioned, and we have expanded upon the attendant difficulties elsewhere.[33] The problem is that the standard recourse with Frege puzzles—that there are distinct modes of presentation—is not at hand; the cases are not just another instance of the evening star and the morning star. The difference is that while thinking of Venus as the morning star is compatible with thinking of Venus as the evening star, thinking of Jews in an anti-semitic way is *in*compatible with thinking of them non-anti-semitically. If one is rational, there is no single state of mind such that one is thinking of Jews both ways; it is either one way or the other. On the other hand, we can think of the planet as the morning star while simultaneously thinking of it as the evening star.

References

Anderson, L. and Lepore, E., 2013, "What Did You Call Me? Slurs as Prohibited Words," *Analytic Philosophy*, 54, 3, 350–363.

Antonelli, G.A. and May, R.C., 2012, "Quantifiers and Determiners," in Russell, G. and Graff-Fara, D., eds., *Routledge Companion to Philosophy of Language*. New York, NY: Routledge, 342–354.

Brontsema, R., 2004, "A Queer Revolution: Reconceptualizing the Debate over Linguistic Reclamation," *Colorado Research in Linguistics*, 17, 1, 1–17.

Camp, E., 2013, "Slurring Perspectives," *Analytic Philosophy*, 54, 3, 330–349.

Chalmers, D., 1996, *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.

Copp, D., 2001, 'Realist-Expressivism: A Neglected Opinion for Moral Realism', *Social Philosophy and Policy*, 18, 1–43.

Frege, G., 1893, *Grundgesetze der Arithmetik*. Jena: Verlag Hermann Pohle. Ebert, P. and Rossberg, M., trs., forthcoming. Oxford: Oxford University Press.

Glanzberg, M., 2006, "Quantifiers," in Lepore, E. and Smith, B., eds., *The Oxford Handbook of Philosophy of Language*. Oxford: Oxford University Press, 794–821.

Hom, C., 2008, "The Semantics of Racial Epithets," *Journal of Philosophy*, 105, 8, 416–440.

Hom, C., 2010, "Pejoratives," *Philosophy Compass*, 5, 164–185.

Hom, C., 2012, "A Puzzle About Pejoratives," *Philosophical Studies*, 159, 3, 383–405.

Hom, C. and May, R., forthcoming, "The Inconsistency of the Identity Thesis," in Lepore, E. and Yi, J., eds., *Language and Value*. Oxford: Oxford University Press.

Horn, L., 1989, *A Natural History of Negation*. Chicago: University of Chicago Press.

Hornsby, J., 2001, "Meaning and Uselessness: How to Think About Derogatory Words," in French, P. and Wettstein, H., eds., *Midwest Studies in Philosophy*, *Volume 25*, *Figurative Language*. Oxford: Blackwell Publishers, 128–141.

Jeshion, R., 2013, "Slurs and Stereotypes," *Analytic Philosophy*, 54, 3, 314–329.

33. See Hom and May, forthcoming.

Kaplan, D., 1999, "What is Meaning? – Explorations in the theory of Meaning as Use," brief version—draft no. 1.

Kratzer, A., 1999, "Beyond Ouch and Oops. How Descriptive and Expressive Meaning Interact," handout for the Cornell Conference on Theories of Context Dependency 1999.

Kripke, S., 1980, *Naming and Necessity*. Cambridge, MA: Harvard University Press.

Lepore, E. and Anderson, L., 2013, "Slurring Words," *Noûs*, 47, 1, 25–48.

Potts, C., 2005, *The Logic of Conventional Implicatures*. Oxford: Oxford University Press.

Richard, M., 2008, *When Truth Gives Out*. Oxford: Oxford University Press.

Salmon, N., 1986, *Frege's Puzzle*. Cambridge: MIT Press.

Schiffer, S., 1987, "The 'Fido'-Fido Theory of Belief," in Tomberlin, J., ed., *Philosophical Perspectives 1: Metaphysics*. Atascadero, CA: Ridgeview Publishing Company, 455–480.

Schlenker, P., 2003, "A Plea for Monsters," *Linguistics and Philosophy*, 26, 1, 29–120.

Soames, S., 1987, 'Direct Reference, Propositional Attitudes, and Semantic Content," *Philosophical Topics*, 15, 47–87.

Whiting, D., 2013, "It's Not What You Said, It's the Way You Said It: Slurs and Conventional Implicatures," *Analytic Philosophy*, 54, 3, 364–377.

Williamson, T., 2009, "Reference, Inference, and the Semantics of Pejoratives," in Almog, J. and Leonardi, P., eds., *The Philosophy of David Kaplan*. Oxford: Oxford University Press, 137–158.