

An Invocational Theory of Slurs[★]

E. McCreedy and C. Davis

¹ Department of English
Aoyama Gakuin University
mccready@cl.aoyama.ac.jp

² Department of Languages and Cultures
University of the Ryukyus
cmdavis@ll.u-ryukyu.ac.jp

Abstract. This paper proposes a new theory of the semantic/pragmatic function of slurs. On this theory, slurs invoke a preexisting complex of social attitudes and background relating to the slurred group. Interpreters then reason about the intention of the speaker in bringing this complex to salience, which accounts for a range of effects produced by slurs in different cases, modeled using default logic axioms about language use. We then show how these calculations of intention do not fully explain the effects that slurs have in a conversation. In addition to standard pragmatic reasoning about speaker intentions, the invocational component of a slur's meaning results in the imposition of a complex of historical facts and attitudes onto the discourse context, and this very act is itself consequential above and beyond what the invocation says about the speaker's communicative intentions. The theory is compared with other theories of slurs; finally, some ways in which this view could extend to other expressions with and without social meanings are indicated.

1 Overview

Words are powerful. Their power is often not immediately visible, but sometimes it is: the use of certain words in certain settings has striking effects, and there are certain words which produce immediate and even shocking results whenever they are used or even mentioned. One class of such words is *slurs*, words that have as one of their functions the denigration of some group of people. The powerful quality of these words and their complex functions have drawn the attention of many people in various domains, not least in recent years in linguistics and philosophy, where much ink has been spilled on their semantics and pragmatics.

* The authors would like to acknowledge the support of JSPS Kiban C Grant #16K02640, which partially supported this research, and to thank Cameron Domenico Kirk-Giannini, Uli Sauerland and Grégoire Winterstein for helpful comments and discussion. Also, it seems appropriate to acknowledge the irony of invoking Lovecraft in a discussion concerning the language of intolerance and racism, given his own attitudes. We hereby express our wish to disavow any unpleasantness this reference may evoke, while acknowledging the potential futility of such intentions for those who find Lovecraft himself problematic.

This paper proposes a new theory of the semantic/pragmatic function of slurs. On this theory, slurs invoke a preexisting complex of social attitudes and background relating to the slurred group. We use the metaphor of a magical invocation for this process: the utterance of the word itself invokes the complex of information associated with it. This sort of invocation has complex pragmatic effects. After the invocation, interpreters must reason about the intention of the speaker in bringing these attitudes to salience. The process itself can be modeled using a combination of default logic axioms about language use and game-theoretic pragmatics; the first of these is discussed in some detail below. The power of slurs, however, goes above and beyond the reasoning they trigger about a speaker's intentions. By summoning a particular complex of facts and attitudes into the discourse context, the slur has effects that are not under the speaker's control, and which in some cases may conflict with those intentions. After introducing our view, we compare the theory with other theories of slurs, mostly within philosophy, where there are several proposals sharing commonalities with our own, though differing in crucial respects. The paper closes with a discussion of some ways in which the view we present could extend to other expressions with social meanings, where an invocational character is common, and without social meanings, where the facts are more subtle.

2 Metaphor

A sorcerer stands atop a high cliff by the sea. He raises his hands and pronounces a single word; a submerged island rises above the waves, covered with cyclopean masonry and dripping sea plants which make the precise angles of the constructions and their outlines indistinct (cf. [16]). The sorcerer has summoned up a city of the past from beneath the sea, where it had heretofore lain invisible. The sorcerer can do so even without knowing every detail of what lies in the city, how it is arranged, or what the consequences of calling it up will be. Utterance of the summoning word is sufficient for the invocation. No one person observing the summoning can see all features of the conjured object. This includes the sorcerer himself.

We suggest that this case is (surprisingly) analogous to the function of slurs. The sorcerer has used a powerful word to call up a hidden, ruined city; slurs, on our view, also bring a preexisting complex of historical facts and constructed attitudes (stereotypes) about the slurred group to attention, in addition to predicating group membership in the manner of a standard nominal. However, as with the obscured nature of the summoned city, it is hard for any one person to discern exactly what those attitudes are, or what the precise historical facts being deployed are. Speakers and interpreters ordinarily don't have full access to the entire complex of historical facts about any slurred group, or to the precise correlates of the stereotypes surrounding them.

This observation partly explains the difficulties linguists and philosophers have had in spelling out the precise function of slurs; lacking full access to the content, it's hard to see exactly how to state it. Within linguistics, this fact has led to a number of analyses which provide 'placeholders' for the slurring content, such as a simple statement that the speaker has a negative attitude toward the slurred group, or the like (e.g. [26, 17, 9]); in some cases, these are intended as full analyses, but in other cases, they are proxy

for a more substantive discussion, usually disposed of with a promissory note. (The first author of this paper is one guilty party in this practice.) It may be that the apparent ineffability of slurs (cf. [27]) is partially explained by the fact that no one person has full knowledge of the complex of facts and attitudes they invoke: it is difficult or impossible to give a full description of something that one lacks access to. Still, those invoking the slur, like the sorcerer, have some purpose in doing so, whether or not they know precisely what it is they are invoking; it is for their hearers to try to understand what that purpose is. But the invocation has effects above, beyond, and in many cases despite whatever intentions the utterer had or is inferred to have had.

3 Application

To see how the metaphor above applies to the use of slurs, we examine four cases in which they are used. The primary factor involved in choosing these cases is the speaker's and hearer's membership vs. nonmembership in the slurred group, selected because this factor plays a primary role in the computation of the intention behind the slur, not to say the degree to which it injures or otherwise affects those present in the situation of use. We thus consider cases in which two social groups are involved: a privileged group, and a group oppressed by the privileged group (e.g. whites and blacks in the USA, ethnic Japanese and ethnic Koreans in Japan, etc). In the following, *S* is the slur, *SG* is the slurred group, and *PG* is the privileged group.

1. *S* is uttered by a member of *PG* to a member of *SG*: **subordination**.
2. *S* is uttered by a member of *PG* to a member of *PG*: **complicity**.
3. *S* is uttered by a member of *SG* to a member of *SG*: **solidarity**.
4. *S* is uttered by a member of *SG* to a member of *PG*: **accusation**.

In each case, a complex of information relating to the slurred group is invoked by the utterer of the slurring word via its semantics. Our strategy now is to consider how the addressee will interpret the use of the slur; in essence, we aim to spell out the pragmatic effects that the invocation has in each case, together with some part of the influence context has on the inferences drawn. We also consider how the invocation itself, while triggering inferences about the speaker's intentions, has effects that are independent of those intentions. Our discussion is, necessarily, programmatic, but we believe it shows the force of the view.

Before turning to the discussion of the individual cases, we want to point out one remaining issue: the problem of defining the privileged group *PG*. Consider, for example, the n-word; here, the privileged group has historically been white people, however defined. What about other groups? What happens if a slur is used by a someone from a group that is in some sense orthogonal to the *PG/SG* dichotomy? Such uses go beyond the scope of this paper, but are clearly worthy of further investigation, which we must leave for another occasion.

We now turn to a consideration of the individual cases.

3.1 Case 1: Subordination

Why would a member of *PG* choose to call up a complex of negative attitudes and historical fact about *SG*, to a member of *SG*? The most obvious reason is to remind

them of their supposed subordinate position; this presumably is the main function of racist terms when used directly to a member of the slurred group. In the absence of further information, this is the likely effect of the use of a slur in this setting. However, one can find contexts in which the intent is different, for instance when a well-meaning member of *PG* says to a *SG* friend about a third, racist, individual *R* the sentence in (1), where the intention involves an attempt to conjure *R*'s attitudes for explicatory purposes.

(1) It's because he thinks of you as a *S*.

The presence of this additional information about the speaker and her intent allows the hearer to infer a different pragmatic function for the slur, in a manner we will address in detail later in the paper.

With this observation, the flexibility of our account is already visible; since the summoned content itself is not attitudinal on the part of the speaker (unlike the various accounts of e.g. [27, 10, 17, 3] i.a.), the speaker doesn't end up committed to any particular attitude, but rather only to the existence of a complex of historical facts and attitudes. It is well-known that theories which both take slurs to be attitudinal and to express their slurring content in an expressive manner have difficulties with projection: given that expressive content is in general independent of truth-conditional content ([26, 27, 17], i.a.), the speaker is expected to be committed to negative attitudes toward *SG* anytime *S* is used, which is not universally the case (cf. [12]). It is nevertheless a persistent intuition that the content of slurs in some sense projects invariably, regardless of what we might infer about the speaker's attitude in a given context. This tension has made it difficult to answer the question of whether or not the expressive content of slurs universally projects. This fraught question is given a natural answer in the invocational theory: projection always occurs, but is nonattitudinal; what projects is simply the complex of historical fact and attitudes associated with the slur (which we will sometimes abbreviate as *C* in what follows). This is in accord with the intuitive linguistic facts.

Our account will allow for the calculation of speaker intent in cases like (1), and in some cases such uses do indeed seem to properly convey the nonhurtful intent of the speaker to her audience. Nevertheless, such uses are fraught, and the risk of failure is high. As anyone familiar with what is likely the most volatile slur in (American) English, the n-word, should be aware, for many members of the target slurred group, use or even mention of the slur by any member of the privileged group is out-of-bounds, regardless of intent.

3.2 Case 2: Complicity.

When a member of *PG* uses *S* to another member of *PG* the complex of content associated with the subordination of *SG* is invoked. Since the way in which that complex is intended to be used is pragmatically derived via reasoning about intention (as we will spell out shortly, following [18]), the functions of doing so are quite various depending on the context. We can separate out two easy cases based on the known attitudes of the speaker. Let's suppose first that the speaker is a known racist; then the use of *S* will ordinarily be processed as an invitation to the hearer to participate in racist attitudes, induced by forcing content associated with them onto the 'conversational scoreboard' ([15, 13]). This is an invitation to be complicit in the project of subordination [28].

Note that, because of the expressive character of the invocational content of slurs, the hearer has little say in determining whether the invoked content is allowed to go down on the permanent record of the discourse (as it were); as Potts [27] and others have stressed, expressive content is inflicted on the hearer, not put up for discussion and possible entry into the common ground.

In the absence of additional context, and if little or nothing is known about the speaker, producing a sense of complicity is the most likely interpretation when slurs are used by one member of *PG* to another member of *PG*. Alternatively, suppose the speaker is known to be a non-racist or is obviously summoning *C* with a non-racist intention. Then the use of *S* will be inferred to have a different (intended) function, often one used to criticize the relevant perspectives, as in the quotative example (1) above, though not always. In example (2), adapted from [3], the recovered content is unavoidably racist, regardless of whatever the speaker's intentions may have been.

(2) Some of my best friends are *S*s.

As in Case 1, reasoning about a speaker's intentions is not always an ameliorative to the use or mention of the slur. A fellow member of the privileged group might take offense despite an accurate calculation of the speaker's innocent intentions. This offense stems, we take it, from the very act of having summoned *C* itself, which a fellow member of the privileged group might feel to be out-of-bounds regardless of what the speaker might have had in mind. This could stem from a number of factors, not all of which are indicative of empathy toward the subordinated group. For example, a white listener might not want to be reminded of the complex of historical facts and attitudes invoked by an utterance of the *n*-word, perhaps out of a desire to avoid the consequences of those facts, emotional and otherwise.

3.3 Case 3: Solidarity.

A perennial problem for theories of slurs that have them semantically expressing negative attitudes about the slurred group is the existence of appropriated uses of slurs, i.e. uses on which the slur is used between members of *SG* with a meaning which is in some sense positive. Theories which posit the explicit introduction of negative attitudes have serious problems with appropriated uses, as they have to introduce ambiguities or underspecification specifically for these cases, which is theoretically unsatisfying. For us, however, these uses are no different than those above (mechanically speaking). The speaker uses *S*, a term which calls up a complex of historical fact and attitude which indicates her subordination and oppression due to membership in *SG*; the hearer is affected in the same way by these facts and attitudes because of his also being part of *SG*. It is part of background knowledge that when people are affected in similar ways by negative situations, a feeling of shared problems and sympathy arises; it is this sympathetic feeling of solidarity that can be inferred from the ingroup use of slurs. Because speakers are aware of this set of facts, the use of *S* can be leveraged and reclaimed as a solidarity marker. Thus, the invocational semantics together with the pragmatics of invocations directly yields the existence of appropriated uses of slurs. We believe that the conventionalized character of many of these uses arises from the usual process of

semantic change apparent in expressive content more generally (cf. [4]). This type of semantic conventionalization can lead to distinct lexical items, as in the case of the two variant pronunciations of the n-word, only one of which is associated with the solidarity use.

We note that, as with other attempts at “deflected” uses of slurs, solidarity uses are not necessarily endorsed by all members of the subordinated group, as is apparent from debate among African Americans about such uses of the n-word. There are, moreover, complexities involving the appropriate contexts for such appropriated uses, in particular, questions regarding whether the appropriated term should be used in “mixed company”. The complexities are illustrated with an anecdote from Randall Kennedy (who is African American):

In my final year at my wonderful high school, St. Albans School for Boys, a black friend jokingly referred to me as a n***er in the presence of one of our white classmates. If he and I had been alone, I might have overlooked his comment or even laughed. But given the presence of the white classmate, I concluded immediately that a show of forceful disapproval was imperative.
[11]

Before closing, we mention one more issue. In the quote above, we have chosen to mask the slur in question, despite the fact that, in the original text, Kennedy chose to write the slur out. In fact, the first word in the title of Kennedy’s book is the very word we have masked, but without the masking. While some readers may roll their eyes at our hypersensitivity to social taboos or even accuse us of magical thinking (an accusation we would welcome, as our main thesis should make clear), we hope to show in what follows that the issue here is not trivial. Neither of the authors of this paper is black. It is not for us to say whether or when black speakers should or mention this word, but we are certainly uncomfortable spelling it out, let alone pronouncing it. As we hope to show, this trepidation stems from respect for the power of invocation that the mere mention of these words can trigger; we should respect them, just as fictional characters in worlds where magic is operative must be careful with dangerous spells. That said, scholarly conventions have led us to cite the title of the book itself in its unmasked form (the reader is hereby warned). Our quandary is, we think, indicative of the fundamentally intractable issues that arise with the question of how to quote these words. We return to this issue at the end of the paper.

3.4 Case 4: Accusation.

The final case we consider is one in which a member of the slurred group *SG* uses *S* to a member of the privileged group *PG*; this case, to our knowledge, is not discussed in the literature. This case is perhaps the most difficult, and the most fraught to consider. It requires the subordinated individual to call up the complex of attitudes and facts that has resulted in her subordination and exhibit it to a member of the group that is oppressing her; this cannot be anything but a painful act, and one which requires internalization and awareness of the damaging circumstances of oppression. For the invoker, it is a reminder of her oppressed status; for the listener, it is a reminder of his status as oppressor, and

indeed an exhibition of the injustice that he participates in. As such, one function of this kind of use is to bring into the light the relative statuses of the discourse participants; the reasons for doing so are various, of course, and the effects will depend in part on the attitudes and sensitivity of the addressee, but often there is an accusatory quality to this use. Consider for example an utterance of (3) in a Case 4 setting; here, the use of *S* feels like a calling out of the addressee, at least at some level.

(3) It's because I'm an *S*, isn't it?

4 Mechanism

Slurs are invocational. They call up a partially knowable complex of facts and attitudes relating to the slurred group: or so we have argued. We believe that the discussion in the previous section shows the flexibility and utility of this view, by exhibiting how the invocational theory applies to a variety of cases of the use of slurs. Now it is time to make it concrete by embedding it in a formal semantic and pragmatic theory. Doing so is the task of this section.

4.1 Semantics

We begin by providing a semantics for slurs. As stated previously, we take them to predicate membership in the slurred group *G* in addition to invoking *C*. The predicate *G* just picks out the slurred group itself, just like the corresponding neutral group term. We take *C* to be expressive content, making our view a descendent of others which claim slurs to have an expressive component to their meanings ([31, 27, 17, 9], i.a.). These two assumptions mean that slurs bear what has been referred to as *mixed content*, as they carry information in both at-issue and expressive domains. Thus, (4) is a schematic lexical entry for a slur. Here, the superscripts on the types indicate the type of content they carry: '*a*' indicates at-issue or ordinary vanilla truth-conditional content, and '*s*' indicates (in this case) expressive content.

(4) $[[S]] = \lambda x.G(x) \blacklozenge C : \langle e, t \rangle^a \times t^s$

Here is a sample derivation with this lexical entry, again schematic.

(5) A is a *S*.

(6)

$$\frac{a : e^a \quad \frac{\lambda x.G(x) \blacklozenge C : \langle e, t \rangle^a \times t^s}{\lambda x.G(x) : \langle e, t \rangle^a \bullet C : t^s}}{\lambda x.G(x) : \langle e, t \rangle^a}}{G(a) : t^a}$$

Interpretation in this system is of whole trees. Any *t*-typed instances of not-at-issue content are extracted from the tree and placed into a set comprising the second element of a two-dimensional meaning; the truth-conditional content is the first element. This mechanism has been somewhat controversial in the literature as it is arguably not fully compositional (e.g. [2]), but, although fully compositional alternatives do exist ([2, 8]),

they don't make much difference for present purposes, so we will put these issues aside here.

(7) Denotation: $\langle G(a) : t^a, \{C : t^s\} \rangle$

A final question about the semantics of slurs involves the invoked content. Why is the complex C summoned by the use of a slur consistently negative in quality? One possible answer would be that non-negative terms for particular groups simply aren't parsed as slurs; but it is rather difficult, or even impossible, to come up with terms of this kind that feel neither slurring nor racist. Our hypothesis is that the negative character of slurring terms is a consequence of the kind of content that is summoned. It is easy to find atrocities in history that are perpetrated by one group of people against another, *qua* groups (e.g. [29] for one case relevant to current concerns): it is much harder to find cases where one group actively does good things for another *qua* group. Moreover, though stereotypes about some group of people might be positive in quality ("Asians are smart", "black people are good at sports"), these stereotypes remain racist and demeaning, and thereby negative (cf. [10]). Since slurs, on our view, simultaneously call up a complex of historical facts and attitudes about a group and predicate membership in that group of an individual, even if the complex is "positive" this would be a kind of reduction of the individual to the complex, which is itself an insult. If this is correct, there is simply no way to use invocations of this kind in a way that has no dehumanizing character, which explains the lack of a positive analogue to slurring terms.

4.2 The Pragmatics of Invocational Purpose

Empirical adequacy requires that we show how the above inferences can be derived; the reader uninterested in such things can safely move to the next section. For this purpose, we adapt the theory of [18], where default logic and signaling games are used to model resolution of emotive underspecification in expressive items, as in (8), in which the expressive adverbial *fucking* can get very different interpretations depending on the extralinguistic context.

- (8) Trump just fucking tweeted about North Korea again!
- a. The speaker is an ardent Republican and red-hat-wearing Trump supporter
 \rightsquigarrow **positive attitude expressed**
 - b. The speaker is an anti-imperialist who often joins antifa demonstrations \rightsquigarrow
 negative attitude expressed

The basic idea is that linguistic agents have default assumptions about language use codified using normality conditionals, and, after observing linguistic behavior, certain of those conditionals will apply. Update of hearer belief states with the derived consequents alters the probability of particular interpretations so that proper resolution can occur in a signaling game setting. In this paper, we will not consider the game-theoretic aspect of the analysis, for two reasons. First, we want to concentrate on the content of slurs, and show how that content can play into hearer beliefs about what a speaker intended to convey; the game-theoretic aspect of interpretation is not really needed for this purpose. Second, unlike the case of emotive adjectives, which present the hearer with

a highly circumscribed choice between two interpretative options (positive/negative), inferences about speaker intention in the use of slurs are completely open-ended, which makes the precise specification of a signaling game extremely difficult and unwieldy at the very minimum, and possibly even a quixotic task of a similar type to the frame problem [6]. We consequently leave it aside here.

The reason to choose default logic for this analysis is its nonmonotonic quality. This makes it useful for modeling real-world inference, which is why it has seen extensive use in artificial intelligence applications (cf. [22]). In nonmonotonic logics, it is possible to override initial conclusions by adding additional information to the set of premises for an inference; as is well-known, this is the usual structure of reasoning about generics [23], and has also proven useful in modeling Gricean reasoning [24, 19]. Inferences associated with resolving speaker intention in particular cases of language use also follow the usual reasoning patterns associated with default logic such as *Specificity* and the *Nixon Diamond*, as extensively argued by [18]. We therefore adopt this perspective here.

The analytical strategy is to produce a set of world knowledge axioms which properly derives the readings which are intuitively present in the cases under consideration. We turn to this task now.

Case 1. This case will give a feel for how the strategy works. The axiom below says that, given that a speaker belongs to *PG* and addresses a member of *SG* with a slur *S*, ordinarily one can conclude that she intends the subordination of her hearer. This axiom, of course, stands or falls with the provision of a more precise analysis of what it means to subordinate someone: candidates include the elimination of their power to choose, a downgrading of their epistemic authority [21], the introduction of a more general kind of epistemic injustice [7], or, most likely, all of the above. We will not choose between these options here, leaving the notion as a primitive for the purposes of this paper. Here and below, '[*S*]' indicates a slur, '[**S**]' a sentence, and '⊆' an inclusion relation on linguistic expressions.

- (9) $PG(s) \wedge SG(h) \wedge Say(s, [S]) \wedge [S] \subseteq [S] > Intend(s, Sub(h))$
 'A *PG*-speaker saying *S* to a *SG*-hearer normally intends the subordination of the hearer.'

The key point is that, given the specificity property of nonmonotonic logics – that a conditional with a more specific premise overrides a less specific one with a conflicting conclusion – the above axiom will fail to apply in certain cases. Consider, for example, the following axiom, which is in conflict with the above.

- (10) $PG(s) \wedge SG(h) \wedge friend(s, h) \wedge Say(s, [S]) \wedge [S] \subseteq [S] \wedge Quote(s, [S])$
 $> \neg Intend(s, Sub(h))$
 'A *PG*-speaker saying *S* to a *SG*-hearer in a quoted context normally does not intend the subordination of the hearer in case the two are friends.'

Since this axiom is more specific than the one above, it will override it, and a speaker trying to (for example) explicate the behavior of some third party will not necessarily be understood as trying to perform an act of subordination. Of course, this effect may arise

anyway, since the effects of invocation are not fully under the control of the speaker; this point will be discussed in more detail later in this section.

Case 2. This case involves use of slurs between members of the privileged group, which, as argued above, generally involves a call to complicity in the project of subordination. We model this using the axiom below. Here, the conversational participants are meant to jointly come to believe in their superiority; the intended notion of belief here is ‘active belief,’ where the proposition in question enters active awareness (cf. [5], i.a.).

- (11) $PG(s) \wedge PG(h) \wedge Say(s, [S]) \wedge [S] \sqsubseteq [S] > Intend(s, Bel(\{s, h\}, Superior(\{s, h\})))$
‘A *PG*-speaker saying *S* to a *PG*-hearer normally intends that the two acknowledge their joint status as privileged or superior individuals.’

Case 3. This is the case of reclaimed slurs used for purposes of solidarity (which is not always the case, of course, as discussed above). The notion of solidarity is explicated in the axiom as a reduction in social distance between the conversational agents between the time of speech *t* and the time *t'* after the utterance is processed; social distance here is construed as a real-valued interval included in $[-1, 1]$, representing a range of values the agents take to be epistemically plausible degrees of distance between them, as in [20], where it is used to model the action of honorifics.

- (12) $SG(s) \wedge SG(h) \wedge Say(s, [S]) \wedge [S] \sqsubseteq [S] > Intend(s, SD_r(s, h) < SD_t(s, h))$
‘A *SG*-speaker saying *S* to a *SG*-hearer normally intends to reinforce social solidarity between the two agents (stated here as a reduction of social distance).’

Case 4. The axiom for this case is very similar to that for Case 2; it also indicates that the speaker’s intent in using *S* in such a context is to bring a proposition into the active awareness of the hearer, again one involving the hearer’s place in some purported social hierarchy. The difference is that in this case, the speaker places herself in an asymmetrical position with respect to the hearer, thereby highlighting their own oppression.

- (13) $SG(s) \wedge PG(h) \wedge Say(s, [S]) \wedge [S] \sqsubseteq [S] > Intend(s, Bel(h, Oppress(h, s)))$
‘A *PG*-speaker saying *S* to a *SG*-hearer normally aims at an accusation of the hearer as participating in oppression.’

These axioms will derive the “default” interpretation of the four cases described. Of course, the reader is warranted in asking where the axioms themselves come from. As with other aspects of world knowledge, the answer goes beyond the domain of linguistics proper. They may be probabilistic generalizations based on repeated exposure to particular uses of the slur in particular contexts, or they may derive from more complex reasoning. In any case, the default logic formalizes both the default interpretation of speaker intentions in the situations we consider, and also allows for violations of those expected interpretations in particular contexts.

We should note, finally, that the above is purely a way for listeners to try to recover what the speaker had in mind when using the invocation. It is not, of course, a description of the actual effects of the slur, which may be quite unpredictable; it is a common observation in the literature (and in real life) that the actual offense drawn from the use of slurs is often unexpected to those who use them. We take this up in the next section.

4.3 Beyond Intentions: The Consequences of Invocation

Returning to our metaphor: the sorcerer on the cliff, in uttering the word of invocation, succeeds in summoning a submerged city. It is *the city itself* that is summoned by the sorcerer. It is not any particular person's view of, experience with, or opinion about that city. The consequences of the summoning on a given third party are fundamentally uncontrollable, although in some contexts relatively predictable, in a way dependent on the relation with and experience of the city itself that the third party involved has. The consequences may well go beyond whatever the sorcerer's intentions were. Although a third party may wonder what those intentions were, and at a magical court martial such questions may help determine a proper sentencing, the effects of the summoning are immune from such questions of intention. The effects of a bullet are not affected by the intentions of the shooter, although those intentions are relevant to the question of whether those effects should result in a verdict of murder or manslaughter.

Imagine that, while reading this paper, you are suddenly hit in the face. Your eyes water. Your vision blurs. You can feel the tissue under your eye begin to swell. Now, you look up to see who or what has hit you. You may discover a number of things. You may find that someone has hit you on purpose, and is preparing to hit you again. Or you may find that someone nearby has carelessly swung their arm into your face. What you learn about the intentions of the person who hit you will likely affect how you interpret the action, and how you react. But whatever your conclusions, your face still hurts.

Similarly, in our own Case 1, we can easily imagine (and some of us have undoubtedly experienced) cases in which a privileged group member has uttered a slur in the presence of a subordinated group member, with no obvious problematic intent. In fact, the *PG* may even have intended to indicate solidarity with *SG*. Such intentions, even when correctly inferred by the listener, are often (or perhaps never) sufficient to eliminate the fundamental impact of the slur, which resides in its power to invoke. Our invocational view of slurs in some sense makes them the ultimate dynamic meaning bearers: they don't merely refer to something language-external, but bring that thing into (proximal) existence. No speaker has access to all facts about the thing being summoned, nor to all perspectives on it. By summoning the thing itself, rather than a particular perspective on the thing, the speaker performs a fundamentally creative, and potentially highly destructive, speech act.

4.4 Beyond Use: Mentions of Slurs and Invocation

One recurring issue that arises in the discussion of slurs is how to appropriately mention them. Returning to our example of the *n*-word, there are many people who find any mention of the word at all to be out-of-bounds, including in quotational contexts. This sense extends to the written form, leading many writers to use instead a euphemism, as we have done here with the *n*-word. This looks, on the surface, no different than avoidance behaviors of other taboo vocabulary. In fact, at least one analysis of slurs essentially reduces their impact to their taboo status: [1], discussed in more detail in the next section.

But our magical metaphor is relevant to this issue as well. In the magical case, words themselves have the power to summon and invoke. At least in some magical cosmologies, mere mention of the word is sufficient to trigger its effect: consider the common

fantasy/horror trope of a hapless individual who runs across a spellbook and begins to read it aloud, thereby calling up unknowable and malign entities and getting himself in serious trouble. If slurs are indeed magic words in the sense we have indicated, it is no surprise that they work the same way. Their invocational force is linked to the word itself, as a linguistic object: in this sense, it's not the content (semantics) of the word that matters in every case, but its mere sound (phonology). Mere mention of the word suffices to invoke *C*; indeed, the word itself need not even be mentioned. In one (in)famous case, an utterance of the word *niggardly* triggered outrage (and counter-outrage), due to the phonology of the word containing a substring sounding like the n-word [25]. Hence taboo avoidance or indulgence in this case is not merely a formal consideration. It amounts to a decision as to whether to trigger invocation.

5 Comparison

As indicated above, our theory differs from most others on the market in not making any reference to attitudes or perspectives of the speaker within the semantics. Its closest analogue is the work of Anderson and Lepore [1], according to whom slurs are simply prohibited words, the use of which is disallowed by social norms. Their view is similar to ours in taking the content of slurs to be in a sense neutral, and so it also avoids the projection problem: since slurs are simply taboo, their use is always prohibited, but commits the speaker to no particular content. It differs from ours in not giving the slurs content beyond taboohood. Their theory does not explain why slurs are offensive, beyond the fact that they are prohibited, but ours does: the historical facts themselves are offenses. No one wants to be reminded of ugly, violent stereotypes and history, but that is precisely what slurs call up from beneath the social waters, where they remain submerged most of the time.

For Elisabeth Camp, slurs involve the introduction of perspectives [3], where a perspective is 'an integrated, intuitive way of cognizing members of the targeted group.' (335). She takes a perspective not to commit the speaker to 'any specific content,' but to motivate 'certain feelings' about the slurred group while not being a feeling itself. Her analogy is to honorific terms such as *tu/vous*, slang expressions, and other thick terms like *cowardly* [30]; we will discuss the former case at the conclusion of this paper. For Camp, these are approximately dispositions to structuring one's thought in a certain way. For the particular case of slurs, the speaker indicates that she holds a perspective which disposes her to characterize members of the slurred group *G* in terms of that membership, which in turn leads to other beliefs. There seems to be little relation between invoking complexes of facts and stereotypes and the semantic introduction of perspectives, so the two pictures are quite distinct, at least initially, though likely there are deep relationships between the two. Still, our view appears to have at least one clear advantage: it immediately explains the existence of appropriated uses of slurs (our Case 3 above) without any need for spurious ambiguities, whereas if slurs introduce perspectives it is necessary to posit a different perspective which the appropriated use can introduce.

Finally, the view of Christopher Hom [10] has certain similarities to our own. According to Hom, slurs are predications of roughly the form '*x* is P_1, \dots, P_n because of

being a member of SG , where the properties are those stereotypically and historically associated with the slurred group. This is also not attitudinal. Hom places this content in the truth-conditional dimension, because he believes that no predication of a slur can be true. He claims that the causal connection indicated by his semantics invariably fails, since no individual has these properties as a result of their group membership, and so the predication always comes out false; but this immediately leads to wrong consequences in cases of negation and other truth-conditional operators, just as with bivalent treatments of presupposition; this is one reason to advocate for an expressive treatment. We leave a more detailed comparison of his idea of social practices with our invoked complexes for future work.

6 Conclusion and Extension

This paper has proposed a view of slurs on which they are magical words, invoking knowledge of historical complexes of oppression. The utility of this view was shown for some specific cases of the use of slurs, and the mechanisms, semantic and pragmatic, by which slurs receive their efficacy were indicated; finally, the view was briefly compared with several related philosophical positions.

There is much left to do on this philosophical and linguistic project. For the case of slurs itself, it remains to explore differences in meaning between particular slurs, as well as the details of particular cases. Also, the formal pragmatic explication we provided uncovered a need to provide a model-theoretic framework in which notions like oppression, subordination, solidarity and awareness can be explored. This is necessary for a full understanding of how slurs work and the structures in which they operate. This theoretical understanding is useful, or even required, for the project of defusing slurs, as one piece of a general project of diminishing the power of black-magical invocations, via an accurate picture of their effects.

We also believe that this general view has a much broader application. In particular, we think it can be straightforwardly extended to other expressions that carry stereotypical social meaning, such as honorifics, which evoke hierarchical relations, and gendered pronouns, which evoke gender roles and stereotypes, not to mention the interplay of the two, a project currently in process [20]. The general invocational view also seems to have ramifications far beyond what is usually taken to be the domain of social meaning. Names have an invocational character, as is well understood in popular culture (e.g. [14]); perhaps even such things as common nouns do as well. The exploration of this magical side of language is a project that we believe can open a new doorway on the formal analysis of meaning in both linguistics and philosophy.

References

1. Anderson, L., Lepore, E.: What did you call me? Slurs as prohibited words. *Analytic Philosophy* 54(3), 350–363 (2013)
2. Barker, C., chieh Shan, C.: Donkey anaphora is in-scope binding. *Semantics and Pragmatics* 1(1), 1–46 (2008)
3. Camp, E.: Slurring perspectives. *Analytic Philosophy* 54, 330–349 (2013)

4. Davis, C., Gutzmann, D.: Use-conditional meaning and the semantics of pragmaticalization. In: Csipak, E., Zeijksstra, H. (eds.) *Proceedings of Sinn und Bedeutung* 19. pp. 197–213 (2015)
5. Fagin, R., Halpern, J.: Belief, awareness, and limited reasoning. *Artificial Intelligence* 34, 39–76 (1988)
6. Fodor, J.: *The Mind Doesn't Work That Way*. MIT Press (2000)
7. Fricker, M.: *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press (2009)
8. Gutzmann, D.: Use-conditional meaning (2015), to appear from Oxford University Press
9. Gutzmann, D., McCready, E.: Quantification with pejoratives (2015), to appear in *Pejoration*, R. Finkbeiner, J. Meibauer, and H. Wiese (eds.): Amsterdam, Philadelphia: John Benjamins (Series "Linguistics Today")
10. Hom, C.: The semantics of racial epithets. *The Journal of Philosophy* 105, 416–440 (2008)
11. Kennedy, R.: *Nigger: The Strange Career of a Troublesome Word*. Knopf Doubleday (2008)
12. Kirk-Giannini, C.: Slurs are imperatives (2017), paper presented at *Integrating Approaches to Social Meaning* workshop at ESSLLI2017
13. Langton, R., West, C.: Scorekeeping in a pornographic language game. *Australasian Journal of Philosophy* 77, 303–319 (1999)
14. Le Guin, U.: *A Wizard of Earthsea*. Parnassus (1968)
15. Lewis, D.: Scorekeeping in a language game. In: Bäuerle, R., Egli, U., von Stechow, A. (eds.) *Semantics from Different Points of View*. Springer Verlag, Berlin (1979)
16. Lovecraft, H.: *The call of Cthulhu*. *Weird Tales* (1928)
17. McCready, E.: Varieties of conventional implicature. *Semantics and Pragmatics* 3, 1–57 (2010)
18. McCready, E.: Emotive equilibria. *Linguistics and Philosophy* 35, 243–283 (2012)
19. McCready, E.: *Reliability in Pragmatics*. Oxford University Press (2015)
20. McCready, E.: *Honorifics and social meaning* (2017), under contract for Oxford University Press
21. McCready, E., Winterstein, G.: Negotiating epistemic authority. In: Kurahashi S. et al (ed.) *New Frontiers in Artificial Intelligence*. Springer, Berlin (2017), to be published
22. Nute, D.: Defeasible logic. In: Gabbay, D., Hogger, C. (eds.) *Handbook of Logic for Artificial Intelligence and Logic Programming*, Vol. III, pp. 353–395. Oxford University Press (1994)
23. Pelletier, J., Asher, N.: Generics and defaults. In: van Benthem, J., ter Meulen, A. (eds.) *Handbook of Logic and Language*, pp. 1125–1177. MIT Press (1997)
24. Perrault, R.: An application of default logic to speech act theory. In: P. R. Cohen, J.M., Pollack, M. (eds.) *Intentions in Communication*, pp. 161–186. MIT Press (1990)
25. Potts, C.: The conventions for expressive content words (blog post on language log) (October 2003), <http://itre.cis.upenn.edu/~myl/languagelog/archives/000035.html>
26. Potts, C.: *The Logic of Conventional Implicatures*. Oxford University Press (2005), revised version of 2003 UCSC dissertation
27. Potts, C.: The expressive dimension. *Theoretical Linguistics* 33, 165–198 (2007)
28. Stanley, J.: *How Propaganda Works*. Princeton University Press (2015)
29. Tirrell, L.: Genocidal language games. In: Maitra, I., McGowan, M. (eds.) *Speech and Harm: Controversies over Free Speech*, pp. 174–221. Oxford (2012)
30. Väyrynen, P.: *The Lewd, the Rude and the Nasty: A Study of Thick Concepts in Ethics*. Oxford (2013)
31. Williamson, T.: Reference, inference and the semantics of pejoratives. In: Almog, J., Leonardi, P. (eds.) *The Philosophy of David Kaplan*, pp. 137–159. Oxford University Press, Oxford (2009)