SLURS ARE VERDICTS

by

Jacob Sapir

A Thesis Submitted in

Partial Fulfillment of the

Requirements for the Degree of

Master of Arts

in Philosophy

at

The University of Wisconsin-Milwaukee

May 2021

ABSTRACT

SLURS ARE VERDICTS

by

Jacob Sapir

The University of Wisconsin-Milwaukee, 2021
Under the Supervision of Professor Peter van Elswyk

Speaker-orientation views (Hom 2008; Neufeld 2019; Camp 2013; Anderson and Lepore 2013; Williamson 2009; Jeshion 2013; Bolinger 2017) explain why slurs are offensive in terms of what slurs reveal about the general point of view of their speakers. Directive theories (Kirk-Giannini 2019) explain slurs are offensive by predicting that slurs issue directives that direct others to adopt their speaker's point of view. As Kirk-Giannini (2019) notes, speaker-orientation views face the problem of old news. Slurs can communicate novel offensive content even after a speaker's general point of view is known. Directive theories, meanwhile, face a novel problem I dub the directive problem. Contrary to what such theories predict, the felicity of a slur does not rest on the felicity of its associated directive. There is thus need for a novel theory of slurs, one which solves both problems. This paper proposes a theory that does that. Slurs communicate ideological verdicts: they represent that disrespectful, ideological practices apply to their targets and reveal the ideological formation their speaker is a part of. Slurs therefore neither issue directives (hence the solution to the directive problem) nor just reveal a general point of view (hence the solution to the problem of old news); rather, they apply that point of view in a context—much like a judge might apply a legal doctrine to a specific case.

TABLE OF CONTENTS

1. Introduction

Slurs are derogatory words like bitch and kike. Theories of slurs explain how slurs function by identifying their core semantic and pragmatic properties.

This paper develops a theory of slurs by initially focusing on two ways that slurs and ideologies (more specifically: ideological formations) interact. First, slurs signal that disrespectful, identity-based ideological practices apply to their targets. Second, slurs are sensitive to ideological context, by which I mean the ideological formations that are contextually salient.[1]

A practice is any pattern of behavior that people rely on in their everyday lives. An example of a practice is a norm, which is a practice that is partly sustained because people value its persistence.[2] A practice is identity-based when it singles out certain identities for special treatment. I leave the notion of an identity-based practice that is disrespectful untheorized. The rough idea is that for an identity-based practice to be disrespectful, it must give those who regard it as legitimate reason to disrespect members of its target identity.[3]

An ideological formation is a network of practices that mediates how a core group of participants understand and interact with outsiders and each other. A practice is thus ideological, if it is part of an ideological formation. I borrow this notion of ideological formation from Haslanger (2017). However, to appropriate terminology from Geuss (1981), I want my analysis of slurs to be descriptive rather than pejorative. Consequently, my usage diverges from Haslanger's in that an ideological formation need not be morally or epistemically deficient. More generally, I'm happy to count

[1] The question of how slurs and ideologies interact has been under-discussed in the literature. Hom (2008), Kukla (2018), and Swanson (forthcoming) are notable exceptions. I discuss Swanson's position on slurs in §2.1, and in fn. 11 in particular. Hom's theory of slurs is criticized in §4. Kukla is discussed in fn. 15 of §5.

[2] I borrow this way of understanding what a norm is from Elster (1989).

[3] For a good analysis of how slurs are disrespectful see Kukla (2018).

words like fascist or pig as slurs, even if it turns out that the ideological formations they are associated with are moral and epistemically in good condition.

An example of an ideological formation is the network of practices that underlies how people interact with each other in a classroom. The network of practices that constitutes the ideological formation of a classroom explains why students raise their hands in order to get their professors' attention, why students sit in chairs rather than on the floor, and why the professor is permitted to interrupt students, but students are not permitted to interrupt the professor.

The contention that slurs signal disrespectful practices apply to their targets is most clear in the case of misogynistic ones that are used to enforce patriarchal norms.[4] Imagine, for example, a cat-caller, who calls the woman he has complimented a misogynistic slur after she ignores him. In doing so, the cat-caller signals that he thinks patriarchal norms apply to the woman, such as, e.g., the norm that women should pay attention to men regardless of whether they want to.

The contention is also supported by reflection on what is communicated by slurs such as kike and fag. Such slurs can be understood as identifying the people they name as the targets of various ideological practices that involve treating Jews and gays, respectively, with disrespect.

Once it is recognized that slurs are connected to ideological practices, it becomes additionally clear that slurs are sensitive to ideological context. Consider, for instance, the slur kike. It is associated with a large number of anti-semitic practices. But those practices are not all part of the same ideological networks. For this reason, it is seldom the case that all practices associated with kike are invoked at once. Instead, what is typically invoked is some proper subset that are associated with a single network of practices—that is, a single ideological formation.

[4] This observation enjoys empirical support. Felmlee et. al. (2020) aggregated together 2.9 million tweets over the course of a week that contained "instances of gendered insults (e.g. 'bitch,' … 'slut' or 'whore')" (25) and found that "hostile, sexist tweets … aim to promote traditional, cultural beliefs about femininity" (16).

After all, consider the following question regarding what it is to be a Jew: If a person's only Jewish parent is their dad, do they count as Jewish? The answer changes depending on the ideological formation of who you ask. An Orthodox Rabbi will answer no. But a Nazi bureaucrat might answer yes. That these membership questions are sensitive to one's ideological formation indicates that slurs are sensitive to ideological formation, too. For suppose Angela calls Bernie a kike. If we then want to understand what Angela is communicating, we need to know what conception of Jewishness is being invoked. Consequently, we need to know whether Angela's conception of Jewishness is like that of the Orthodox rabbi, or like that of the Nazi bureaucrat.

I name the theory this paper develops verdict literalism. I divide the theory into two components, each of which impose conditions on what it takes for a word to count as a slur. The first component, the semantic component, is the conjunction of the claims (1), (2), and (3).

1. Each slur is associated with a collection of disrespectful identity-based, ideological practices.
2. Each predicative use of a slur invokes those and only those practices that are associated with both the slur and some referenced or represented ideological formation.
3. For any slur S, S is correctly predicated of a person X, just in case it correctly represents X as being such that at least one of the practices S invokes applies to them.

The second component, the hermeneutic component, is the conjunction of (4), (5) and (6).

4. Each slur is associated with a niche set of ideological formations, each of which has a practice of using the slur in question.
5. When a speaker uses a slur, the ideological formations that use the slur are sufficiently sparse that audiences typically can recognize the speaker as representing one of the ideological formations that has a practice of using the slur.
6. A speaker who represents an ideological formation while using a slur regards the ideological formation's practices as legitimate, and—unless they are making

> reference to an ideological formation that they are not representing—invokes with their slur just those practices that are associated with both the slur and the ideological formation they represent.[5]

The semantic component deals with what slurs overtly communicate. It holds that a word is a slur only if it is (a) associated with disrespectful identity-based ideological practices, and (b) accurately represents a target if and only if at least one of those practices applies to them.

The hermeneutic component deals with how audiences interpret speakers. It holds that a word is a slur only if it leads audiences to identify its speaker as a member of one of the ideological formations that have a practice of using the slur. The underlying idea is that slurs are words which only a sparse number of ideological formations have a practice of using, such that any instance of slur-use provides audiences with enough information that they can identify the formation the slur-user is representing.[6]

I name this view verdict literalism, because one of its key claims is that slurs overtly communicate ideological verdicts. They communicate verdicts, in that they communicate that practices apply to their targets; and they do this via their overt semantics, in that they do not communicate their verdicts via such mechanisms as conventional or conversational implicature. This means verdict literalism is at odds with views that identify a slur's meaning as equivalent to the meaning of its neutral counterpart. Instead, verdict literalism holds that word-pairs like kike and jew or bitch and woman differ from each other when it comes to their literal meaning.[7]

---

[5] The *unless* clause is meant to account for cases where a speaker makes reference to an ideological formation that is known not to be their own. For example, suppose a Jew says, "Harvard used to treat *Jews* as *kikes*." Here the speaker is understood to be invoking the practices of an ideological formation (the one belonging to the Harvard of years past) that they do not represent.

[6] Verdict literalism is thus similar to views due to Anderson and Lepore (2013) and Bolinger (2017), in that it holds that slurs are words that only a select number of people have a practice of using. On my view, however, slurs do not need to be taboo, and though the facts about who uses a slur can reveal a speaker's ideological allegiance, the slur itself primarily communicates information via a conventional mechanism.

[7] Views that endorse such an equivalence between slurs and their neutral counterparts include Williamson (2009), Jeshion (2013), Anderson and Lepore (2013), and Kirk-Giannini (2019).

My defense of verdict literalism is abductive. The theory can explain the same kind of data that its competitors can (§2), and it also helps explain ideological change (§3). At the same time, it solves the problem of old news, which rules out many of its rivals (Hom 2008; Anderson and Lepore 2013; Camp 2013; Bolinger 2017; Neufeld 2019) (§4) and also solves a novel problem faced by a theory defended by Kirk-Giannini (2019) (§5).[8]

2. The offensive profile of slurs

This section considers the following data. First, slurs can communicate offensive content (§2.1). Next, slurs can communicate offensive content even when negated (§2.2) and even when their speaker attempts to disavow any implied offensive content (§2.3).

2.1. Slurs communicate offensive content

Consider again the case of the cat-caller who calls the woman who has ignored him a misogynistic slur. The cat-caller's slur-use communicates offensive content. Hence the datum. A good analysis of this datum both predicts that the datum is true and clarifies the nature of the content communicated. To this end, verdict literalism predicts that slurs can be offensive by communicating three types of offensive content. First, slurs communicate that their speakers regard illegitimate practices as legitimate. Second, they communicate ideological threats. Third, slurs communicate the conditions under which members of their speakers' ideological formations find communicating such ideological threats acceptable.

(i) To start, verdict literalism predicts that not every use of a slur communicates that illegitimate practices are legitimate. For one thing, not all slurs are associated with illegitimate practices — this is plausibly the case with the slur fascist. For another, speakers can use slurs without thereby endorsing the disrespectful practices those slurs

[8] Kirk-Giannini discusses a problem similar to the one I identify in a footnote. Nevertheless, the problem Kirk-Giannini addresses is distinct from the one that I have come up with.

invoke. This is plausibly the case when a Jew says the Germans treated Jews like kikes or when a lesbian calls herself a dyke. On a natural reading of both cases, the Jew is not endorsing the Germans' treatment of Jewish people, and the lesbian is not suggesting that it is right that she be disrespected.

The existence of these exceptions is consistent with verdict literalism's tenets. Per the semantic component, slurs always communicate that disrespectful identity-based practices apply to the people they're predicated of. But to say that a disrespectful practice applies to someone is not yet to communicate that the practice is legitimate. Instead, slurs communicate that disrespectful practices are legitimate by first revealing their speakers' ideological formations. When a speaker reveals that their ideological formation includes practices that their slur invokes, and when those practices are themselves not legitimate, only then does the speaker reveal that they regard illegitimate practices as legitimate.

How then does the use of a slur reveal a speaker's ideological formation? The basic idea is this: slurs are only used by specific ideological formations. So, the use of a slur provides audiences with strong evidence that the slur-user belongs to one the of the formations associated with the slur. Additionally, audiences have antecedent beliefs about the formations a given speaker is plausibly a part of. Consequently, a slur's capacity to reveal the allegiance of its speaker may be modeled as follows. Let S be the set of ideological formations that have a practice of using a slur and let K be the set of ideological formations that an audience thinks a speaker is plausibly a member of. A slur reveals a speaker's ideological allegiance to an ideology i when i is at the intersection of S and K.

Thus, consider again the case of the Jew who says

7. The Germans treated Jews like kikes.

Such a speaker is not plausibly a member of an anti-semitic formation. And the slur itself is not only used by anti-semites. So the intersection of S and K yields the result that the speaker is not anti-semitic. Additionally, the hermeneutic component holds

that a speaker invokes the practices associated with both their ideological formation and their slur unless they are making reference to an ideological formation that is not their own. In this case, the speaker is clearly referencing the ideological formation of a group of Germans. So, the speaker is claiming that the referenced German ideological formation included anti-semitic practices.

Next, consider the case of the lesbian who calls herself a homophobic slur. An audience who knows that the woman in question is a lesbian will not antecedently believe that the woman could be plausibly homophobic. Additionally, though used by formations that are homophobic, is also used by ideological formations that regard being the target of homophobic practices as a source of esteem. Consequently, in calling herself a dyke, the woman reveals herself to be a member of this latter type of ideological formation. In calling herself the slur, therefore, the woman communicates (i) that homophobic practices apply to her and (ii) that she thinks this makes her worthy of esteem.

Lastly, consider the cat-caller. The cat-caller's use of a misogynistic slur reveals he is allied to one of the ideological formations that uses the slur. Some of those formations are misogynistic, but others are not. The woman he is cat-calling, however, has a sense of her local ideological topography, and so she knows that the non misogynistic ideological formations with a practice of using bitch do not have a practice of yelling at women they see on the street. Consequently, the cat-caller's slur-use reveals he regards practices that are disrespectful to women as legitimate. But such practices are not legitimate. Hence the cat-caller has communicated the first type of offensive content.

(ii) The next type of offensive content verdict literalism predicts slurs can communicate are what can be understood as ideological threats. It is one thing for a speaker to reveal the ideological formation they are a member of. It is another for the speaker to additionally communicate that their formation's disrespectful practices apply to you. To learn a speaker is part of a bad ideological formation is to learn something that is merely offensive. But to learn that they think some of the disrespectful practices they regard as legitimate apply to you is to learn something that is also threatening. For

example, in cases where being the target of the disrespectful practices is a matter of choice, the additional information provides one with an incentive to change one's behavior in order to no longer be those practices' target.

Thus, consider again the cat-caller. Per the above discussion, the cat-caller's slur-use reveals that he regards patriarchal practices as legitimate. But, per the semantic component, the cat-caller's slur-use also reveals that he thinks some of those practices apply to the woman who ignored him. He thus communicates a kind of ideological threat, of the form: disrespectful practices apply to you, and I regard those disrespectful practices as legitimate. Moreover, the ideological threat puts pressure on his target to change how she behaves. It gives her an incentive to respond to her future cat-callers, lest she be subject to the their (potentially violent) disrespect.[9]

(iii) The last content type is this: slurs can communicate—via a conversational implicature—that the speaker regards their issuing of an ideological threat contextually acceptable. Note that in this case I mean acceptable in a variety of senses, of which two are significant. First, I mean acceptable as in personally acceptable: that is acceptable from the speaker's own point of view. Second, I mean acceptable as in ideologically acceptable: that is, acceptable for members of the speaker's ideological formation.

One reason to think that slurs communicate this third type of content is that it helps explain how slurs can in effect colonize a context-type. The use of a homophobic slur against a gay couple walking in a public park, for example, can discourage future gay couples from making use of that public space. After all: it makes sense for a gay couple to avoid a public space if they think members of a given ideological formation find issuing ideological threats in that public space acceptable. So if, as is predicted by verdict literalism, members of my ideological formation find issuing ideological threats

[9] Manne (2017, 2020) observes that misogyny, unlike, say, antisemitism, is more geared towards the *management* rather than elimination of the identity-group it targets. Misogynistic slurs are thus more likely to yield these sorts of pragmatic reasons to change one's behavior than other types of slurs: From the point of view of the antisemite, there is no such thing as the right kind of Jew. But from the point of view of the misogynist, there *is* such a thing as the right kind of woman.

in this public space acceptable is a message that slurs can communicate, it follows that the use of a slur in a public space can have precisely such a colonization effect.

Though there is this kind of abductive evidence to believe that slurs can communicate this sort of acceptability content, the main motivation for the thought that slurs implicate that their use is contextually acceptable is a principle I borrow from Swanson (forthcoming).

> Swanson's Principle:
> Speakers generally conversationally implicate that what and how they communicate is contextually acceptable.

The motivation for Swanson's principle may be illustrated with a case of patronizing over-explanation.[10]

> Alyosha and Ivan are on a road trip. Ivan is driving and Alyosha is navigating. Alyosha is mad at Ivan for not liking enough of his social media posts. As retribution, Alyosha provides Ivan instructions in a patronizing way. For example, he frequently asks Ivan supplements his directions with rhetorical questions such as the following: "You remember what stop signs do, don't you? Those are the red octagons you see by intersections sometimes. You remember how many sides an octagon has, right?"

Alyosha's style of explanation is clearly patronizing. Swanson's principle explains why. Speakers are assumed to speak only in a style they find acceptable. Yet Alyosha's communication style is usually reserved for children who don't know the rules of the road. As a result, Alyosha generates the conversational implicature that it is acceptable for him to communicate with Ivan in a style reserved for kids. Moreover, he generates this implicature for a variety of senses of acceptable. He communicates that it is personally acceptable, but he also communicates that it is acceptable for anyone in a relevantly similar position: that is, for anyone who likewise is mad at their friend for social-media related injustices.

[10] I owe the idea of illustrating Swanson's principle with a case of patronizing over-explanation to Zachary Ferguson.

Thus, suppose a speaker uses a slur to issue an ideological threat. Given Swanson's principle, such a speaker will also issue an acceptability implicature to the effect that communicating the ideological threat is acceptable, given the context. First, the speaker will communicate that they find the ideological threat personally acceptable; and second, the speaker will communicate that others like them find the ideological threat personally acceptable, too. And in the context of an ideological threat, others like them will just be other members of their ideological formation. The third type of offensive communicated content is thus predicted by Swanson's principle, together with verdict literalism earlier prediction that slurs communicate ideological threats.[11]

2.2. Slurs can remain offensive even when negated

So far, I have discussed the datum that slurs communicate offensive content. Next, I discuss the datum that a speaker's use of a slur can remain offensive even when the slur is embedded under negation.

For example, the following sentence comes from Neufeld (2019).

8. He's not a kike, he's a muslim.

Alternatively, consider this dialogue.

9. Coach: Alright kid, let me ask you, are you a bitch?
10. Buratino: No I'm not, sir.
11. Coach: Then why don't you act like it!

Both of these examples illustrate that negating a slur does not fully cancel its problematic implications. First, it is reasonable to think that someone who utters (8) is

[11] Swanson (forthcoming) applies his principle to the case of slurs in a similar way that I do. However, there are two major differences. First, Swanson assumes that slurs and ideologies are connected without explaining how, whereas verdict literalism explains that the connection is semantic. Second, Swanson assumes that the revelation of ideological membership is enough to cue certain ideological practices, whereas verdict literalism does not depend on this assumption, since it predicts slurs invoke ideological practices through what they overtly communicate.

anti-semitic, even though the claim that the slur kike applies to someone is being directly denied. Next, Buratino's denial that a misogynistic slur applies to him still seems to render him complicit in Coach's misogyny.[12]

Verdict literalism's explanation for the negation data is simple. Speakers who use slurs pejoratively provide their audience with evidence that they think the disrespectful practices with which the slur is associated are legitimate. Moreover, per Swanson's principle, they then conversationally implicate that providing such evidence and invoking such practices is contextually acceptable. Since neither of these two messages is transferred by way of the slur's literal meaning, neither is targeted by a speaker's negation of the slur's semantic content.

Thus, consider again example (8). Per verdict literalism, a speaker who utters (8) simply denies that there exist any anti-semitic practices that apply to the person they are talking about. But, through their use of kike, they still supply evidence of their allegiance to an antisemitic ideological formation and hence they still implicate that the practices they have invoked are legitimate.

Next, consider the dialogue. Coach's question implicates that certain patriarchal and misogynistic practices are legitimate. When Buratino denies that such practices apply to him personally, Buratino fails to deny the implication that those practices are nevertheless legitimate. Hence the impression of complicity.

2.3. Slurs are hard to cancel

An additional nuance in the offense-data is this: the offensive content communicated by slurs is often hard to cancel.

12. Look, I'm not anti-semitic, but Sasha is a kike.
13. I know it is not acceptable for me to call you this, but why are you being a bitch?

[12] The first to have discovered this datum is Potts (2007). Potts' explanation involves the claim that slurs conventionally implicate that their speakers accept derogatory content, the idea being that conventional implicatures project out of negation. Others who discuss this datum include Williamson (2009), Camp (2013), and Kirk-Giannini (2019).

In (12) the speaker is attempting to claim that they are not anti-semitic. Yet this is hard to believe, given that they called someone an anti-semitic slur in the very same sentence. Meanwhile, in (13), the speaker attempts to disavow the belief that their use of bitch is acceptable in the context. But this is yet again hard to believe, given their subsequent use of the word in that very same context.

The challenge, then, is this. Per the discussion in (§2.1), Slurs can reveal their speaker's ideological formation and issue the acceptability implicature that they think their slur-use is acceptable. Both of these messages, moreover, are communicated via non-semantic means. Yet discourses (11) and (12) reveal that they are difficult to cancel. How come?

Verdict literalism's response to this challenge is as follows. First, since speakers who use slurs indicate their ideological allegiance by supplying their audience with evidence to that effect, the failure to cancel such an implication is just the result of the evidence supplied by the slur being sufficiently strong that it swamps the contrary evidence supplied by the speaker's attempt at a disavowal. Second, acceptability implicatures are in general hard to cancel, because it is hard to do so without admitting to practical irrationality, and audiences are often more likely to interpret a speaker's admission that they are irrational as untruthful than they are to regard the speaker as genuinely irrational. After all: only someone who is practically irrational could truthfully suggest that doing X is unacceptable, immediately after doing X.

Thus, here is the explanation for (12). The evidence provided by the speaker's use of kike swamps the contrary evidence provided by the speaker's contention that they are not anti-semitic. Hence, the speaker's disavowal is just read as either a lie or as an instance of false consciousness.

Next, here is the explanation for (13). Discourse (13) features the speaker putatively admitting that they know their use of bitch is unacceptable, even though they immediately thereafter make use of the word. It thus involves the speaker seemingly making the admission that they are irrational. For if they really believed that it is unacceptable for them to use the slur bitch, and if, additionally, they were rational, they

would simply not use the slur in the first place. But they did. Since audiences are sooner to decide that the speaker is lying and behaving rationally than that the speaker is truthful but behaving irrationally, the speaker of (13)'s attempt to disavow their acceptability implicature is likely to be unsuccessful.[13]

3. Slurs and ideological change

So far I've discussed how verdict literalism explains the offensive profile of slurs. Next I demonstrate that verdict literalism predicts a way slurs can contribute to ideological change.

The basic idea is this. Practices are not directly visible. So a person who is committed to a particular practice figures out what specifically they are committed to, not by observing the practice they have a commitment to directly, but by observing how people who take themselves to be following the practice conduct their lives. In particular, such a person figures out what they are committed to by paying attention to how people use slurs, which—per verdict literalism—can be used to communicate whether certain practices are applicable or not.

The use of a slur thus has the capacity not just to invoke certain ideological practices, but also to change a group's understanding of what those practices consist in—especially when the applicability-conditions of those practices are vague.

Thus, consider the following toy example. In a high school, Claudia cheats on her boyfriend whom she is going to break up with anyway. The high school includes a practice of shunning those who are excessively promiscuous, but it's vague whether Claudia's behavior counts. In such a case, we can imagine that the boyfriend decides that it does, and that he therefore decides to invoke the applicability of the practice by calling Claudia a slut to anyone who will listen. Finally, we can imagine that this shifts

[13] Since the mechanism that explains why acceptability implicatures are hard to cancel has to do with the difficulty of plausibly professing practical irrationality, a class of cases where an acceptability implicature *can* be cancelled is when a speaker convincingly claims they have issued their slur in a non-rational way, such as via a Freudian slip.

public opinion and thus the content of the practice. Where once it was vague whether Claudia's behavior counts as excessively promiscuous, now, because of the boyfriend's actions, it is certain. Hence a case where a slur not only invokes certain practices, but in fact contributes to ideological change by developing the content of the very practices invoked.

4. The problem of old news

I finish up this paper by demonstrating that verdict literalism dialectically outcompetes its rivals. To start, I discuss the problem of old news. First, I demonstrate that the problem rules out a large number of extant rivals. Second, I demonstrate that the problem can be solved by verdict literalism.

The problem of old news affects a class of theories, which I call speaker-orientation views.[14] Speaker orientation views (Anderson and Lepore (2013); Camp (2013); Jeshion 2013a; Bolinger 2017) explain why a slur is offensive in terms of what the speaker's slur-use reveals about either (a) the speaker's willingness to break taboos or (b) the speaker's general beliefs, desires, or cognitive dispositions, but which make no mention of (c) the speaker's willingness to communicate ideological threats.

An example is a view proposed by Anderson and Lepore (2013). Anderson and Lepore argue that the distinctive offensive profile of slurs is wholly explained by the fact that slurs are "prohibited words." They thus explain the offensive profile of slurs simply by referencing what the use of a slur reveals about a speaker's willingness to break taboos.

Another example is a view proposed by Camp (2013). Camp explains the offensive profile of slurs by holding that slurs "conventionally signal allegiance" to what she calls

[14] The name speaker orientation views was originally given to the class by Kirk-Giannini (2019). However, I've modified Kirk-Giannini's definition so that the class of speaker orientation views better and more explicitly includes all and only those views that are negatively affected by the problem of old news.

"derogating perspectives", or special cognitive dispositions to attend to and explain information in particular ways (331).

The problem of old news is this. Suppose a speaker has already indicated—without using a slur—that they both have a misogynistic point of view and that they are willing to break taboos against indicating as much. If so, then the use of a misogynistic slur is predicted by a speaker-orientation views to be old news: the slur should not be able to communicate novel offensive content, since it is precisely the violation of a taboo or the demonstration of the speaker's perspective that makes the speaker's slur-use offensive in the first place. But in fact this prediction is false. A misogynistic slur in such a situation can still convey novel offensive content.

To illustrate, I present a modified version of an example given by Kirk-Giannini (2019, 6).

> It is well known that there is a secluded community called Andropolis that includes all and only misogynists. Specifically, these misogynists believes woman are disposed to have certain negative stereotypical properties and additionally believe that woman deserve the injustices that they in fact face. Consequently telling people you are from Andropolis is taboo. One day, Albert and Benjamin are at a dinner party. Benjamin wants to know whether anyone at the party has moved house recently. For he wants advice on how to choose a moving company. So he asks Albert, "has anyone you know moved house recently?" Albert answers in one of the following two ways, while pointing at a fellow female partygoer named Doris.

14. I actually recently moved to Andropolis, and, oh, that woman over there recently moved to Chicago.
15. I actually recently moved to Andropolis, and, oh, that bitch over there recently moved to Chicago.

Since mentioning that you are from Andropolis is taboo, both (14) and (15) signal that Albert is willing to break anti-misogynistic taboos. Moreover, since Andropolis includes all and only misogynists, both (14) and (15) signal that Albert has a misogynistic point of view. Consequently, speaker orientation views predict that (14) and (15) ought to convey the same information: it shouldn't matter that in (15) Albert uses a slur, whereas in (14) Albert does not. But this is not so. The addition of a slur in (15) seems to convey

active hostility towards Doris that goes beyond the prior revelation of Albert's dispositions and point of view. It is as if whereas previously Albert simply communicated that he was fine breaking taboos to tell people he is a misogynist, now Albert is communicating that he finds it additionally acceptable for him to enforce his misogynistic point of view in the context of the party.

A large number of theories are ruled out by this counterexample. I have already mentioned the proposals of Anderson and Lepore (2013) and Camp (2013). But other views are affected as well. For example, it affects what might be called property views: views which claim that a speaker's slur-attribution is true of someone only if the target of their slur either has or is disposed to have some cluster of negative stereotypical properties. Hom (2008), for instance, holds that misogynistic slurs are offensive because they reveal their speaker believes women have certain stereotypical properties and so deserve to be subject to the injustice they in fact face. And Neufeld (2019) holds that misogynistic slurs reveal their speakers believe women are disposed to possess certain negative properties. But such content is already communicated by Albert when he reveals he is from Andropolis. Consequently, even property views have trouble with the problem of old news.

More generally, the problem of old news is a problem for any view that sees the use of a slur as an act that does nothing more than signal allegiance to a point of view. For if all a slur does is signal allegiance to a point of view, then an I-belong-to-such-and-such-community sentence can be generated that sends the same message, but which is nonetheless weaker in terms of the content it communicates.

Hence why verdict literalism solves the problem. Though slurs do signal a speaker's allegiance to an ideological formation, they also communicate specific ideological threats: they signal that specific practices are considered applicable in the context. And this is something that an I-belong-to-such-and-such-community sentence cannot communicate.

Thus, consider again the case above. When Albert merely signals he is a misogynist in (14), he does not thereby issue an ideological threat. Therefore, though he implicates

that he finds it ideologically acceptable to reveal he is a misogynist, he does not implicate that he finds it ideologically acceptable to issue an ideological threat that targets the behavior of Doris. Rather, this additional message is only communicated in (15), when Albert uses a misogynistic slur to derogate his fellow partygoer. Hence verdict literalism's solution to the problem of old news.

5. The directive problem

Verdict literalism is not the only theory that can solve the problem of old news. The problem can also be solved by a class of theories I call directive theories. Directive theories, however, face what I call the directive problem, a problem that verdict literalism can solve with ease.

Directive theories hold that slurs direct others to adopt the point of view of their speaker via a conventional mechanism. An example of a directive theory is proposed by Kirk-Giannini (2019). Kirk-Giannini agrees with Camp (2013) that slurs are associated with derogating perspectives. However, Kirk-Giannini thinks that rather than principally signaling allegiance to such perspectives, slurs in fact direct others to adopt them, such that if Allen calls Bernie a kike, Allen is directing Bernie to adopt an antisemitic perspective. This in turn helps Kirk-Giannini solve the problem of old news in a way Camp cannot. For even if people already know that Albert is a misogynist, the additional news that Albert wants them to be misogynists as well is novel content.[15]

Though directive theories solve the problem of old news, they nevertheless face the directive problem. Suppose Allen believes Bernie is incapable of adopting a derogatory point of view towards members of Bernie's own social group. If so, supposing Allen has

[15] A potential example of a directive theory is defended by Kukla (2018). Kukla holds that speech acts which involve the derogatory use of a slur have an illocutionary profile in which, roughly, the target of the speech act is directed to see themselves as they are seen by the speech act's issuer. Depending on how their theory is implemented semantically, it thus has the potential to be directive theory. This need not be the case, however. Verdict literalism is not a directive view, but it can arguably implement Kukla's thesis, too—the idea being that the ideological threats issued by slurs generally have the effect of encouraging others to internalize the point of view of the slur-issuer.

the right psychology, it can be odd for Allen to direct Bernie to adopt such a view. Consequently, if slurs issue directives to adopt points of view, then in such scenarios where the associated directive is infelicitous it should likewise be infelicitous for Allen to call Bernie a slur. Yet this is false. Even if a speaker believes that their target is incapable of adopting a problematic point of view towards members of their own social group, and hence would never attempt to direct their target to think as much, a speaker can still without any infelicity call their target a slur. Hence the directive problem: slurs may be used without infelicity even in cases where the issuing of the associated directive to adopt the point of view is infelicitous.

The following case illustrates the objection.

> The year is 1944 and Roger Chesteron has been captured by Nazis. Roger accepts British wartime propaganda and so subscribes to a Manichean worldview on which Germans are incapable of redemption or changing their minds about the merits of fascist ideology. The Nazis take him to an interrogation room. Roger prepared for this outcome back in London. He places a cyanide capsule underneath his tongue. As his interrogator asks his first question, Roger exclaims (16) and swallows the pill.
>
> 16. Go to hell, you kraut!

In the scenario above, Roger's use of the slur kraut is felicitous. Yet it plausibly would be odd for him to utter one of the following.

17. Please understand: you are a kraut!
18. Listen to me, I need you to get it through your head that you are despicable!

Both (17) and (18) are explicit directives. If it is odd for Roger to say them, we should expect that any utterance that communicates the same kind of content will be odd to say, too. Hence, if we grant that slurs issue directives, it should be odd for Ivan to say (16). Yet this is not so. Hence a misprediction on the part of directive views.

It might be objected that this argument wrongly assumes that if explicit directives are odd to issue, it should likewise be odd to use a slur with the same directive content. Instead, it might be thought that there is a relevant difference: according to Kirk-

Giannini (2019), slurs issue directives implicitly, via a mechanism of conventional implicature, whereas the directives like those in (17) and (18) do so explicitly, as part of their overt semantics.

This however is not so. Suppose for contradiction that whether a message is communicated explicitly or implicitly affects whether it is odd to communicate. Further suppose that fixing refrigerators gives me so much joy, that it would be odd for me to utter (19).

19. My refrigerator stopped working again—I hate this!

If so, then we would expect the utterance of (20)—which only conventionally implicates as much—to not sound as odd.

20. My damn refrigerator stopped working again!

Yet in worlds where (19) is odd to say, (20) is odd to say, too. Hence the idea that whether a message is communicated implicitly makes a relevant difference for whether it is odd to communicate is false.

Kirk-Giannini (2019 fn. 24) considers a problem similar to the directive problem. He points out that it can be felicitous for a speaker to issue a directive even in cases where the speaker does not believe their target will follow the directive's instructions. Kirk-Giannini uses an example featuring Hank, a fictional homophobe, to illustrate. Suppose Hank sees two men kissing at a nightclub and cries out Shame!. Shame! is plausibly a directive that instructs the couple to feel shame. And it is equally plausible that Hank does not think that the couple he is targeting is capable of doing so. Nonetheless, Hank yelling "shame!" at a gay couple isn't psychologically infelicitous. So here it seems we have a case where issuing a directive is psychologically compatible with a belief that the directive will be unheeded.

Call the claim that directives can be psychologically compatible with a belief that the directive will be unheeded the directive compatibility thesis. Though I admit that the directive compatibility thesis is plausible, my objection is sound even if it is true.

To see this, it is important to distinguish the objection I am making from the objection Kirk-Giannini is responding to in the footnote. Kirk-Giannini is responding to an objection whose key premises are (22) and (23).

21. If slurs conventionally issue P-associated directives (directives to adopt some point of view P), then whenever it is infelicitous to issue a P-associated directive, it is also infelicitous to issue a P-associated slur.
22. A belief that a target is incapable of adopting P, does not make it infelicitous to use a P-associated slur.
23. A belief that a target is incapable of adopting P does make it infelicitous to use a P-associated directive.
24. So, there are scenarios where it is infelicitous to issue a P-associated directive but not infelicitous to issue a P-associated slur. (22, 23)
25. So, slurs do not conventionally issue P-associated directives. (21, 24)

And since the directive compatibility thesis implies the falsity of (23), Kirk-Giannini is correct in holding that this objection is unsound. However, the directive problem I have identified is instead based on (24), directly.

24. There are scenarios where it is infelicitous to issue a P-associated directive but not infelicitous to issue a P-associated slur.

And, as I will now argue, even though (23) is false, (24) is true.

To see this, first consider the goal directed activity (GDA) compatibility thesis, a generalized version of the directive compatibility thesis. The GDA compatibility thesis holds that given an activity directed at a goal, it can still be permissible to engage in the activity, even if one believes that the goal is unattainable. Like the directive compatibility thesis, the GDA compatibility thesis is plausible. The goal of playing a soccer game is to win the match. But even if you know that the strength of the team you are going to be up against makes achieving this goal impossible, it can still make sense for you to play them in a match. For instance, perhaps you have made a prior commitment to play the game.

But—and this is the key point—this need not always be the case: there are times when belief in imminent failure is enough reason to make participation no longer make sense. Consider, for example, a student who decides against asking a professor for a letter of recommendation, because he knows the professor will decline. In such a case, the student's knowledge of imminent failure makes it no longer make sense for him to ask his professor for a letter.

Let's now reconsider the example above. Let us stipulate that it doesn't make sense for the officer to utter (17). And then let's check whether it can still make sense for him to utter (16).

16. Go to hell, you kraut!
17. Please understand: you are a kraut!

Given that we have stipulated that uttering (17) does not make sense for Roger to say, it must be that for him the belief—that his captors are incapable of complying with his demand—is sufficient to render his uttering (17) pointless. Nevertheless, it is intuitive that for such an officer uttering (16) still makes sense. After all, his utterance of (16) does not seem to be susceptible to the same kind of failure as his utterance of (17). Hence, again, the directive problem: the infelicity of the officer uttering (17) does not rule out the felicity of the officer uttering (16), and so directive views are wrong to predict that slurs are felicitous just when their associated directives are felicitous.

I conclude by noting that verdict literalism solves the problem quite simply. Per the view, speakers who use slurs do not issue directives. So when Ivan calls his captors krauts, he is only communicating that he thinks the anti-German practices of his era apply to his captors and further that he regards those practices legitimate. In short, since in (16) Roger is not making any demands of his captors, he can coherently view his uttering (17) as pointless and yet nonetheless view (16) as a sensible way to express disrespect.

6. Conclusion

Verdict literalism thus dialectically outcompetes its rivals by solving both the problem of old news and the directive problem. As a way of wrapping up, I now apply the theory to a real-life example.

On July 20, 2020, congressman Ted Yoho called his colleague Alexandria Ocasio-Cortez (AOC) a "fucking bitch."[16] The context was as follows. AOC had recently posited a connection between rising poverty and rising crime rates. Yoho took issue with this and decided to confront AOC on the steps of the Capitol. Specifically, he started by telling AOC that "you are out of your freaking mind." AOC responded that Yoho was "rude." Finally, Yoho reacted by calling AOC a misogynistic slur.

It is hard to know with certainty what was going through Yoho's head during the incident. But from the point of view of a reasonable observer, Yoho's use of a slur can easily be understood as expressing his belief that AOC's status as a prominent progressive congresswoman, coupled with her calling him rude, were improper—that her gender, progressive politics, and refusal to be deferential to her male colleague made her deserving of disrespect. In the language of verdict literalism, he was expressing that certain disrespectful, identity-based practices, perhaps representable as a collective practice of treating those who violate the rule (R) with disrespect, applied to her.

(R) Women whose political opinions are incorrect are not permitted to be politically outspoken.

Naturally such behavior by the congressman was misogynistic and wrong. Verdict literalism helps explain the details. Instead of claiming that Yoho achieved a misogynistic effect simply through the breaking of a taboo or by issuing of a directive, verdict literalism explains that he was invoking specific practices. Instead of claiming that Yoho communicated offensive content solely through what he revealed about his

[16] <https://thehill.com/homenews/house/508259-ocaasio-cortez-accosted-by-gop-lawmaker-over-remarks-that-kind-of>.

own point of view, verdict literalism explains that he also communicated offensive content by issuing an ideological threat. And instead of claiming that therefore Yoho must have been committed to the view that his target has or is disposed to have some cluster of negative properties, verdict literalism instead explains that his slur-use could have just as easily been motivated simply out of a desire that certain patriarchal practices be maintained.

References

Anderson, Luvell, and Ernie Lepore. 2013. “Slurring Words.” Noûs 47 (1): 25–48.

Bolinger, Renée Jorgensen. 2017. “The Pragmatics of Slurs.” Noûs 51 (3). Wiley-Blackwell: 439–62.

Camp, Elisabeth. 2013. “Slurring Perspectives.” Analytic Philosophy 54 (3): 330–49.

Elster, Jon. 1989. Nuts and Bolts for the Social Sciences. Cambridge University Press.

Felmlee, Diane, Paulina Inara Rodis, and Amy Zhang. 2020. “Sexist Slurs: Reinforcing Feminine Stereotypes Online.” Sex Roles 83 (1): 16–28.

Geuss, Raymond. 1981. The Idea of a Critical Theory: Habermas and the Frankfurt School. Cambridge University Press.

Haslanger, Sally. 2017. “Racism, Ideology, and Social Movements.” Res Philosophica 94 (1): 1–22.

Hom, Christopher. 2008. “The Semantics of Racial Epithets.” Journal of Philosophy 105 (8): 416–40.

Jeshion, Robin. 2013. “Expressivism and the Offensiveness of Slurs.” Philosophical Perspectives 27 (1): 231–59.

Kirk-Giannini, Cameron Domenico. 2019. “Slurs Are Directives.” Philosophers’ Imprint 19: 1–28.

Kukla, Rebecca. 2018. “Slurs, Interpellation, and Ideology.” Southern Journal of Philosophy 56 (S1): 7–32.

Manne, Kate. 2017. Down Girl: The Logic of Misogyny. Oxford University Press.

———. 2020. Entitled: How Male Privilege Hurts Women. New York: Crown.

Neufeld, Eleonore. 2019. “An Essentialist Theory of the Meaning of Slurs.” Philosophers’ Imprint 19.

Potts, Christopher. 2007. “The Expressive Dimension.” Theoretical Linguistics 33 (2): 165–98.

Swanson, Eric. Forthcoming. “Slurs and Ideologies.” In Analyzing Ideology: Rethinking the Concept, edited by Robin Celikates, Sally Haslanger, and Jason Stanley. New York: Oxford University Press.

Williamson, Timothy. 2009. “Reference, Inference and the Semantics of Pejoratives.” In The Philosophy of David Kaplan, edited by Joseph Almog and Paolo Leonardi, 137–59. Oxford University Press.