

Dog Whistles, Covertly Coded Speech, and the Practices that Enable Them

Anne Quaranto
University of Texas at Austin
annequaranto@utexas.edu

Abstract

Dog whistling—speech that seems ordinary but sends a hidden, often derogatory message to a subset of the audience—is troubling not just for our political ideals, but also for our theories of communication. On the one hand, it seems possible to dog whistle unintentionally, merely by uttering certain expressions. On the other hand, the intention is typically assumed or even inferred from the act, and perhaps for good reason, for dog whistles seem misleading by design, not just by chance. In this paper, I argue that, to understand when and why it’s possible to dog-whistle unintentionally (and indeed, intentionally), we’ll need to recognize the structure of our linguistic practices. For dog whistles and for covertly coded speech more generally, this structure is a pair of practices, one shared by all competent speakers and the other known only to some, but deployable in the same contexts. In trying to identify these enabling conditions, we’ll discover what existing theories of communicated content overlook by focusing on particular utterances in isolation, or on individual speakers’ mental states. The remedy, I argue, lies in attending to the ways in which what is said is shaped by the temporally extended, socio-politically structured linguistic practices that utterances instantiate.

1. Introduction

Recently, not only news media but also philosophers of language and scholars from other disciplines have taken up the topic of dog whistles—expressions sending a signal pitched too high for some to hear.¹ The technique itself, however, is much older than the label “dog whistle.”

¹ Though the philosophical literature specifically concerned with dog-whistle speech is recent, the topic has already received rather diverse treatments. Focusing on liberal democratic norms of public discourse, Jason Stanley (2015) describes dog-whistle speech as a propagandistic mechanism which serves to surreptitiously erode those norms, even while appearing to conform to them. Justin Khoo (2017), on the other hand, considers whether dog whistles’ in-group messages are a semantic or pragmatic phenomenon, and argues that they are pragmatic results of hearers’

For example, phrases like “law and order,” “inner city,” and “states’ rights” have long been identified as covert racist codes by political scientists and commentators.² A significant question that hasn’t been adequately addressed by this literature is whether dog whistling requires the speaker’s *intending* to dog whistle. This question is significant both for our political practices and for our theories of meaning and communication. In practice, intention is at issue in much of the public discourse on dog whistles: while critics of language like “law and order” often infer an intention to send a racist message from a speaker’s use of the phrase, or even conflate having such an intention with the mere uttering of the phrase, those accused of dog whistling standardly rebut the accusation by citing their lack of such an intention. As for our theories of meaning and communication, dog-whistle speech is a useful test case, because, despite the obvious differences between dog whistles and the cooperative, sincere assertions that are typically the focus of our theorizing, there are some key ways in which the effects produced by dog-whistle speech are closer to ordinary truth-conditional meaning, than they are to, say, a prestige accent: both dog-whistle effects and truth-conditional meanings present or frame the world a certain way, recur

inferences. Jennifer Saul (2018) distinguishes between dog whistles which the in-group recognize as such and dog whistles which raise certain attitudes to salience without the awareness of those affected. Saul theorizes the latter as a type of perlocutionary speech act which succeeds only when the audience is not aware of the speaker’s intention to dog whistle. The linguists Robert Henderson and Elin McCready (2019) incorporate aspects of social meaning into Khoo’s inferentialist treatment. While political scientists have studied racial codes since at least the early 1980s, the work that’s proved most influential for this recent philosophical literature has been Tali Mendelberg (2001), who argues that dog whistles arise when contradictory social norms have been internalized. The legal scholar Ian Haney López (2014) offers an indispensable historical understanding of dog whistles, tracing their use back to the post–Civil War Reconstruction era.

² For example, during the 1968 campaign season, the Fair Campaign Practices Committee described “‘code words’ with racial implications” as their biggest concern, pointing to examples such as “law and order” and “crime in the streets.” United Press International, “New Racial Smears Held More Subtle,” *New York Times*, Sept. 18, 1968, 22.

systematically across utterances in contexts of the same type, and acquire their significance through repeated association of expression and message over time.³

The question of whether dog whistling requires the speaker's *intending* to dog whistle also doesn't have an obvious answer. On the one hand, it seems mistaken to attribute "sneaky" speech acts like dog-whistling to speakers who didn't intend to deceive or manipulate their hearers, just as "lying" seems inapt to describe a sincere but mistaken assertion. On the other hand, expressions like "law and order," whose use as a racist dog whistle is well established, have been observed to send certain messages even without the speaker's intent or knowledge. Speakers might dog-whistle unintentionally, for instance, when they naively repeat talking points couched in language like "law and order," without necessarily knowing or recognizing the race-baiting effects of their words.⁴ What's most crucial (and least obvious) is which expressions could do so, and when and why.

In this paper, I start from this observation that at least sometimes, speakers unintentionally dog whistle merely by uttering certain expressions, and I explain when and why such unintentional dog-whistling is possible. I propose that dog-whistling unintentionally (and indeed, intentionally) is possible only when there's a certain *structure of linguistic practices* in place—namely, a pair of practices, one shared by all competent speakers and the other known

³ I'll say more about this notion of *message* in section 3, but for now it'll suffice to say that utterances sending the same message make the same kind of contribution to a discourse or serve the same discursive role, which might center on information, attitudes and emotions, or value judgements.

⁴ Another kind of example is social scientists studying such language, who have successfully produced such effects by asking survey participants questions about, for example, "inner city criminals." At least in certain experimental designs, these effects can be produced by speakers who do not intend to communicate racist messages—for instance, in designs where the survey questions are administered by research assistants who are purposefully not fully informed about the effects that the study is designed to measure, so as to avoid introducing researcher bias. For more examples, see Haney López (2014), Hurwitz and Peffley (2005), Mendelberg (2001), and Smith (1987).

only to some, the latter sending a different message but involving the same expression and performable in the same contexts.

Moreover, I argue that *only* by taking this kind of practice-focused approach can we see dog-whistle speech clearly. Identifying the conditions in which a speaker could dog whistle unintentionally requires looking not only beyond the speaker's mind, but even beyond the particular utterance in question, for these enabling conditions are to be found in the structures of and relations between the temporally extended linguistic practices of the speaker's and hearers' shared linguistic community. This focus on communities and their linguistic practices thus reverses the relative priorities, as compared to theories that account for communicated content primarily in terms of individuals' beliefs, preferences, intentions, and other mental states. A practice-focused approach, on the other hand, de-emphasizes the speaker and their communicative intentions, and focuses instead on the patterns of behavior of the linguistic community as a whole and of the sub-communities within it.⁵ Relatedly, focusing on communities involves considering utterances not in isolation, nor merely in their immediate context, but as performances of pre-existing practices, as acts embedded both in a history of usage and in a social and political history.⁶

Although the focus on linguistic communities and their practices is not entirely foreign to analytic philosophy of language, it's often taken to be incompatible with theories of meaning and communication in which facts about individual speakers' or hearers' minds are central.⁷

⁵ Here and below, I use "communicative intention" as a shorthand for "the message a speaker intends their audience to receive," and not in the technical sense of Grice (1957), where the communicative intention is characteristically transparent to hearers, and mutually recognizable as such.

⁶ I'll discuss a number of theorists who take such a practice-focused approach in section 4.1.

⁷ For instance, Kripke's (1982) influential interpretation of Wittgenstein's rule-following paradox posits a communal, pragmatist account of meaning *in place of* the individualist accounts (whether in terms of mental states or behavioral dispositions) that are found untenable.

However, this needn't be the case, for we can think of the practice-lens as one tool among many that are useful for studying linguistic phenomena. My goal in this paper is thus not only to identify the conditions that make unintentional dog-whistling possible, but also to advocate a certain methodology. In particular, a practice-focused approach provides more insight into phenomena for which purely individualist approaches offer only superficial accounts. The case study taken up here thus both motivates and models incorporating the focus on practices into a greater variety of approaches in the philosophy of language, and even into seemingly incompatible individualist approaches.

As I see it, these explanatory and methodological goals go hand-in-hand: It's precisely in attempting to explain how unintentional dog-whistling is possible, that the utility of a practice-focused approach becomes clear. At the same time, from within such a practice-focused approach, dog-whistle speech no longer appears as an anomaly or corruption of the mechanisms enabling ordinary communication, but rather as the product of a perfectly ordinary phenomenon familiar to any speaker of a natural language—heterogenous linguistic communities whose speech practices are shared in varying degrees, as with local dialects or technical jargon. Thus, from the practice focus, dog-whistle speech appears no less relevant to theories of meaning and communication than sincere, cooperative assertion.

1.1 Plan of the paper

To explain how unintentional dog-whistling is possible, I'll first need to describe dog-whistling itself. But since the category of dog whistles is neither clear-cut nor uncontested, I begin in section 2 by identifying a broader category, *covertly coded speech*, which encompasses at least the most paradigmatic dog-whistling, as well as other kinds of speech which I'll argue

resemble dog-whistling in three key ways: First, they're *coded*, meaning they send more than one distinct message to different subsets of the audience.⁸ Second, they're *covertly* coded, meaning that one of these messages is received by all hearers, while the other is received only by some and unnoticed by others.⁹ Third, they're not merely nonce codes, but *recur* across multiple utterances by different speakers on different occasions. By design, the resulting category is somewhat broader than the one assumed in the public discourse and scholarly literature on dog whistles; while "dog whistle" is most often applied to speech whose coded message is racist or derogatory, this is neither required nor presupposed by my calling something "covertly coded speech."¹⁰

With this category of covertly coded speech as the target, I then develop my answer to the central question: Under what conditions could a speaker unintentionally produce such speech? I answer this question in section 3.1, where I identify the enabling conditions of covertly coded speech in terms of a certain structure of linguistic practices. The type of speech described in section 2 is only possible, I argue, when there are *two* distinct practices which involve the same expression and can be performed in (at least some of the same) contexts, but which send different messages, and only one of which is shared by all competent users of the expression. So, it turns out that in certain contexts, an utterance is covertly coded because there is such a pair of practices, and not because the speaker intended to speak in covert code. In section 3.2, I motivate

⁸ Note that this definition of "coded" may diverge from our ordinary use of the term.

⁹ As I'll clarify in more detail below, this category of "covertly coded speech" is distinct from what Saul (2018) calls "covert dogwhistling," which is a sub-type of dog-whistle speech which goes unnoticed even by those whose attitudes are successfully affected or changed by hearing the dog whistle.

¹⁰ Constructing the category of covertly coded speech thus serves to foreground the linguistic and communicative aspects of dog-whistle speech, at the cost of backgrounding (at least in the span of this paper) its social, political, and moral aspects. Although I think these aspects deserve careful philosophical study, I won't be able to do them justice in the space of this work.

this practice-focused approach to covertly coded speech by arguing that, even for *intentionally* coded utterances, explaining the three distinctive features of covertly coded speech will require looking beyond the individual speaker or the particular utterance, and instead to the linguistic community and the structure of its practices.

The cases considered in section 3.2 thus motivate the basic tenet of the practice-focused approach—namely, the explanatory priority of linguistic communities and their practices, rather than, say, individual speakers' mental states or abstract formal relations. Section 4 elaborates on this practice-focused approach to language and communication. In section 4.1, I detail the practice-focused approach more thoroughly, identifying the ways in which it's been used by philosophers and linguists in order to elucidate, not just particular kinds of speech, but also foundational questions about the nature of meaning and reference. In section 4.2 I turn to the notion of *social practice* itself, and describe the intuitive, pre-theoretic notion of practices with which practitioners like ourselves are equipped. Although fully delineating the metaphysics of social practices is outside the scope of this paper, this pre-theoretic notion, I argue, is sufficiently rich for the purposes of a practice-focused account of covertly coded speech.

Finally, in the last section of the paper I consider the place of a practice-focused approach among other approaches to the study of communication. Though the notion of linguistic *practices*—and the attendant community-level, historical facts—de-emphasizes individual speakers' communicative intentions, in section 5 I argue that individualistic or mentalistic approaches of communicated content can and should incorporate practice-structure conditions into their accounts of covertly coded speech. I thus conclude by taking stock of the methodological upshots of this study. As demonstrated by the case of unintentional dog-whistling, the practice-focused approach I advocate in this paper both reveals and explains

aspects of our actual speech practices that are omitted or misunderstood when we think of linguistic communication primarily in terms of particular utterances and individual speakers' or hearers' private mental states. At the same time, theorizing communication in terms of communal, embodied, historical practices is not, in fact, incompatible with other, more individualistic or mentalistic approaches. Rather, the practice-focus provides another lens through which to examine some of our most central data, and it prompts us to re-conceptualize that data in fruitful ways, and thereby to complicate and deepen our technical notions of meaning and communicated content.

2. The Target Phenomenon: Covertly Coded Speech

Though speakers disagree about what is a dog whistle, most opinions converge on a few examples. Among linguistic expressions, “states’ rights,” “welfare,” “inner city,” and “law order” are the most commonly agreed-upon examples. And there also seems to be broad, if implicit, agreement on a few central features: a dog whistle sounds like ordinary speech, but it also conveys something problematic (most commonly, something racist). However, some canonical examples diverge even from this minimal description. For example, the 1988 presidential campaign ad titled “Weekend Passes” is typically considered a dog whistle. The ad purportedly criticized Democratic presidential candidate Michael Dukakis’s position on a prison furlough program, but it also described violent crimes while presenting audiences with a mugshot of Willie Horton, a Black man. Critics charged that the ad worked to stoke audiences’ fears by playing on racist stereotypes which present Black men as prone to violent crime and

sexual violence.¹¹ The ad thus seems to have produced effects similar to those of dog-whistle expressions like “law and order,” conveying a racist message without explicitly mentioning race. Moreover, the ad relied on audiences’ familiarity with many of the same practices that enable “law and order” to function as a racist dog whistle—namely, racist, derogatory, and dehumanizing stereotypes about Black Americans.¹² So, though I mostly focus on linguistic expressions in this paper, this example suggests that pre-existing natural-language expressions are not the only medium for dog-whistle speech.¹³ Indeed, some critics have found dog whistles in public figures’ silence on an issue,¹⁴ and even in their clothing choice.¹⁵ The ordinary notion of dog whistles is further complicated by the fact that some of the central features of dog whistles are not in fact unique to dog whistles. As but one example (we’ll soon see more), the phrase “friend of Dorothy” has been used as an in-group shibboleth by gay speakers, allowing them to safely self-identify as gay only to others in the know, while conveying something innocuous to those not in the know.

¹¹ For a thorough account of the ad’s creation and reception, see Mendelberg (2001), 135–165. Saul (2018, pp. 365–66) analyzes the Horton ad as one instance of what she calls a covert dogwhistle—i.e., a dogwhistle which influences some hearers’ reasoning or behavior without their conscious awareness. Note that my use of “covert” differs from Saul’s. I say more about this difference in section 4 below.

¹² As an anonymous reviewer rightly pointed out, it’s often more difficult to identify the relevant enabling practices for something like a TV ad, than for utterances of linguistic expressions. Though this paper focuses on linguistic expressions, the central proposal in principle extends to non-linguistic cases too. I’ll give the practice-focused account of this example in section 3.2 below.

¹³ For an illuminating discussion of an even wider variety of linguistic and visual dog whistles, see Saul and Drainville (forthcoming), who draw examples from advertising, art, memes, and hand signals, to name but a few.

¹⁴ “The Guardian view on Rashford, Sancho and Saka: let down by dog whistles from Downing Street,” *The Guardian*, July 12, 2021, <https://www.theguardian.com/commentisfree/2021/jul/12/the-guardian-view-on-rashford-sancho-and-saka-let-down-by-dog-whistles-from-downing-street>. Thanks to Ian Proops for bringing this example to my attention.

¹⁵ Marci Robin, “The Internet Is Comparing Melania Trump’s Outfit to a Famous Fictional Nazi,” *Allure*, Oct. 8, 2018, <https://www.allure.com/story/melania-trump-africa-outfit-looks-like-nazi-sympathizer>.

My goal is not to capture this heterogenous ordinary use of “dog whistle,” nor to argue for a revision to that use. Instead, in this paper I seek to further our understanding of dog-whistle speech by identifying the mechanisms responsible for its most distinctive—and most insidious—features. These features are shared not only by the canonical examples of dog whistles (like “inner city” and “law and order”), but also by other kinds of coded speech. So, this paper will focus on a broader category of speech which I call *covertly coded speech*. Covertly coded speech is demarcated by three distinctive features—*differential effects*, *hiddenness*, and *recurrence*. As I’ll argue below, theorizing this broader category of covertly coded speech (rather than derogatory dog whistles in particular) allows us to identify something fundamental about how such speech works—namely, *a general mechanism that stems from the structure of our linguistic practices*.¹⁶

The three features named above—*differential effects*, *hiddenness*, and *recurrence*—help to distinguish covertly coded speech from other similar phenomena. Covertly coded speech is a subtype of *coded speech*, and like all coded speech—e.g., spy codes, Pig Latin, slang, technical jargon—there’s something it communicates to only a select few “in the know.” This is the sense in which coded speech produces *differential effects*: it sends a coded message to a subset of the audience, a message which some competent speakers do not receive. Unlike overtly coded speech like spy codes or Pig Latin, however, covertly coded speech doesn’t seem coded to those who aren’t “in the know”; to them it seems like any other ordinary speech. This is the sense in which the codedness of covertly coded speech is covert or *hidden*: it’s not merely that the coded message is hidden from some hearers, but rather the fact that it’s coded is hidden, and hearers

¹⁶ Granted, this may not be the only way that a speaker could produce the kinds of effects that dog-whistle speech produces. But this is not unique to dog whistles or covert codes; our linguistic practices equip and even over-equip us in a variety of ways—consider, for example, the variety of synonyms or sentence structures a speaker might use to convey one and the same message.

who aren't in the know are unaware that any coded message has been sent. Though the comparison with spy codes and Pig Latin might draw our attention to coded speech which informs, the coded "message," as I use the term, encompasses a broader set of discursive effects than just truth-conditional content. For instance, the coded message could be an affective or evaluative framing, rather than some piece of information.¹⁷ I'll return to this notion of "message" in section 3.2 below, where I'll give a fuller definition in practice-theoretic terms.

So far, we've carved out this subclass of acts which are not only *coded*—i.e., which communicate something to only a select few "in the know"—but moreover *covertly* coded—i.e., whose coded nature goes unnoticed by those not "in the know." Both overtly and covertly coded speech produce differential effects on their audiences; that is, they both send a coded message that only some hearers receive. But overtly and covertly coded speech differ in their effects on those who aren't "in the know." For example, utterances using Pig Latin (e.g., "Isthay isyay earlyclay odedcay") or Caesar cypher (e.g., "UIJT JT DMFBSMZ DPEFE") are overtly coded.¹⁸ Utterances like these don't send a message to those unfamiliar with the conventions of Pig Latin or Caesar cypher; they might recognize that the utterance is meaningful for other hearers, but to them it's gibberish.

In contrast, covertly coded speech sends another message, in addition to the coded one. I call this the *ordinary message* because it's one that's sent by the ordinary, commonly known usage of the expression; it's one that *any* competent user of the expression would also have

¹⁷ That said, affective or evaluative framings could also be understood as semantic, even propositional, contents. Stanley (2015, pp. 125-77) weighs several ways we might extend existing semantic models to this broader variety of discursive effects.

¹⁸ Translation: "This is clearly coded." In Pig Latin, initial consonant sounds are moved to the end of each word, and words beginning with vowel sounds are not changed. "Ay" is then appended to each word. In Caesar cypher or alphabet-shift encryption, each letter is replaced by one that is further down in alphabetical order by a set number of spaces. The expression "UIJT JT DMFBSMZ DPEFE" is encrypted by a shift of 1; i.e., by substituting each letter by its immediate follower, so that B replaces A, C replaces B, etc.

received in the same context, regardless of whether they're also familiar with the expression's narrower usage. Sending this additional, ordinary message is what effectively "hides" the codedness of covertly coded speech, and thus what distinguishes covertly coded speech from overtly coded speech. This ordinary message is received by everyone in the audience—that is, by in- and out-group both. Going forward, I'll use "in-group" to refer to the subset of the audience who receive both the ordinary and coded messages, and "out-group" to refer to the subset of the audience who receive only the ordinary message. So as I use the term, being in the "in-group" only requires familiarity with or competence in the narrow practice; it does not require that the hearer endorses or condones the message, nor does it require that the hearer in fact believes, acts on, or reasons on the basis of the coded message.

To illustrate this distinction between overtly and covertly coded speech, contrast the covertly coded expression "law and order" with an overtly coded utterance in Pig Latin (e.g., "Isthay isyay earlyclay odedcay"). For example, Donald Trump, styled himself "the law and order candidate" during the 2016 presidential campaign,¹⁹ and in the summer of 2020 used calls for "law and order" to express disapproval of the nationwide protests against police brutality against Black Americans.²⁰ Audience responses to Trump's use of "law and order" have made it clear that, at least to some hearers, the phrase communicates racist stereotypes that portray Black

¹⁹ Dan Roberts and Ben Jacobs, "Donald Trump Proclaims Himself 'Law and Order' Candidate at Republican Convention," *Guardian* (July 22, 2016), <https://www.theguardian.com/us-news/2016/jul/21/donald-trump-republican-national-convention-speech>.

²⁰ David Smith, "Trump reaches for Nixon playbook after protests that have rocked America," *Guardian*, June 7, 2020, <https://www.theguardian.com/us-news/2020/jun/07/donald-trump-re-election-nixon-protests-strategy>.

Americans as criminals.²¹ Nevertheless, other hearers have maintained that the phrase only communicates something like “a public sphere free of violence and unrest.” Indeed, the structure of covertly coded speech makes it (at least logically) possible that such claims are sincere. In contrast, *overtly coded speech* like “Isthay isyay earlyclay odedcay” doesn’t send an ordinary message alongside the coded message, and as a result is recognized as coded even by those who can’t “decode” the utterance. Competent speakers would not—and could not—mistake such an utterance for ordinary American English.

Despite my focus on examples from politicians’ speech, covertly coded speech occurs in many other settings. Indeed, it needn’t even be harmful or derogatory, unlike the speech typically labeled a dog whistle.²² For example, in some U.S. hospitals, a PA announcement paging “Dr. Strong” is a covertly coded request for security to assist with a combative patient at the specified location (York & MacAlister 2015, p. 497). As with any covertly coded speech, many who hear such an announcement won’t know *what* this is a code for, nor even *that* it’s a code at all. Perhaps surprisingly, uttering a technical term like philosophers’ “valid” may also send a covertly coded message, even unintentionally, in contexts where some audience members are

²¹ Based on data from the 2016 American National Election Studies Time Series Survey, Drakulich et al. (2020, p. 392) found evidence of Trump’s law and order rhetoric serving as a dog whistle: whereas among the entire set of those surveyed “feelings of warmth toward the police were not significantly related to vote choice,” among those *high* in racial resentment, support for the police *did* affect their choice between Donald Trump and Hillary Clinton. Drakulich et al. (p. 394) conclude that, as during the Civil Rights Movement of the 1960’s, during the 2016 election, “‘support for the police’ seemed to be a signal that mattered particularly to voters with high levels of racial resentment. This finding ... shows that a subset of Americans may be cloaking their concerns about the racial order behind a superficially nonracial support of the police.” See also Drakulich et al. (p. 375) for a summary of findings from six decades’ worth of existing literature and survey data on politicians’ use of “law and order” as a coded racial appeal.

²² Another example of non-derogatory covertly coded speech is the phrase “wonder-working power,” used by George W. Bush in his 2003 State of the Union address. For empirical data of the differential effects produced by this phrase, see Albertson (2015). For further discussion of this and other non-derogatory dog whistles, see Saul (2018), as well as Saul and Drainville (forthcoming), who offer an illuminating discussion of a wide variety of examples, going well beyond just the most commonly cited racist dog whistles like “law and order.”

unfamiliar with the technical use of the term, and where the technical and the ordinary usage are each deployable.²³

But probably the most familiar form of covertly coded speech in our day-to-day lives is double entendre, which sends both a humorous message to those “in” on the joke, and another, contextually relevant message to those too young or too naive to be “in” on the joke (as anyone who’s watched Mae West films as a child can attest). Note that you don’t have to find the joke funny, in order to get what the joke is, to be “in” on it in a minimal sense. And the coded message sent by an utterance of “Paging Dr. Strong” may be received by hearers besides the speaker’s intended audience, especially now that the reader is “in on” this coded practice. In general, hearers can receive a coded message without endorsing it, without sharing the speaker’s sensibilities, and without wishing to sustain the practice of sending that message. Rather, all that’s required is that the hearer is familiar with the relevant linguistic practice (a point to which I’ll return below).

The third distinctive feature of covertly coded speech is that it *recurs*, unlike speech which produces hidden differential effects via a one-off mechanism. The coded message sent by a dog whistle, Pig Latin, or slang is sent *whenever* the expression is uttered in a certain context and before a certain kind of audience. Within a context-type, utterances of, for example, “law and order” send the same coded message, even though they involve different speakers,

²³ As an anonymous reviewer has pointed out, there’s also an important difference between covert codes like “valid” and ones like “law and order”: namely, why the narrow practices would arise at all. The coded use of “law and order” arises, not just in a certain structure of practices, but also in a certain type of sociohistorical and normative contexts—i.e., the practical background in which the coded practices arose and in which they’ve been deployed. There are certain questions about covertly coded speech which we can’t answer, without taking these sociohistorical and normative aspects and contexts into account—questions concerning the practical, social, and political goals and ends of using covertly coded speech. While these are certainly worthwhile questions, I must defer addressing them until future work, as doing so would require shifting our focus to different features of practices and actions, than the ones needed to present and defend the thesis of this paper.

audiences, locations, and times. In contrast, the pragmatically enriched content of a highly context-specific insinuation or conversational implicature might also produce hidden differential effects, but not in a way that's repeatable by other speakers, to other audiences, on a variety of other occasions. For example, a hint might rely on the target hearer's reasoning about the speaker's intentions, on the basis of their shared background or what they know about the speaker's beliefs and desires. Another hearer, who doesn't share the same background or know the relevant details about the speaker, would either fail to receive any message from the hinting utterance or receive a different, misleading one.²⁴

Together, these three features demarcate the category of covertly coded speech—*differential effects*, *hiddenness*, and *recurrence*. However, they don't always appear together, and taken separately, they may appear with other types of speech. As we've already seen, overt codes like Pig Latin produce differential effects, but the fact that they're codes is not hidden from the out-group. Other kinds of speech can produce differential effects, and even hidden ones, via a one-off mechanism that only works on a particular occasion. But when they appear together, these three features indicate the category of *covertly coded speech*. In this paper, I focus only on speech which exhibits all three of these features. Hence this paper will not explain, for instance, how dog whistles are first coined, nor does it encompass utterances which produce differential effects but not in a systematically recurring way. But this doesn't mean that the category of covertly coded speech is therefore merely ad hoc or gerrymandered. To explain why, in the next section, I turn to the source of the three distinctive features of covertly coded speech, and argue that they're not merely the result of some coincidental similarity among hearers. Rather, a given utterance of some covertly coded expression sends a certain coded message *because* such

²⁴ Indeed, this is the mechanism proposed in Clark (1992), which I discuss in section 5.1 below.

utterances reliably send that message—that is, because there exists a certain shared practice which the given utterance instantiates. And uttering a covertly coded expression (e.g., the dog whistle “law and order”) can produce covertly coded speech because there exists a particular pair of practices (one narrow, one widespread) which I’ll describe in more detail in the next section.

3. A Practice-Focused Account of Covertly Coded Speech

As I’ve constructed the category in the previous section, covertly coded speech involves expressions which, when they’re uttered in a certain type of context, systematically send two messages at once: a coded message to the in-group, and an ordinary one, as per the commonly known usage, to all hearers. Now that the target phenomenon is more clearly in our sights, I turn to identifying the conditions in which a speaker can produce covertly coded speech. I propose that a speaker’s communicative intention is not necessary for covertly coded speech, nor is it always sufficient. Rather, what makes covertly coded speech possible is a pair of practices which involve the same expression and which can both be performed in at least some of the same contexts, but which differ in that one practice is shared by all competent users of the expressions, and the other is familiar to only some. When, moreover, an audience includes some who are familiar with both of these practices and others who are familiar with only the widely shared one, a speaker uttering the expression for which there are two such practices will produce covertly coded speech, whether intentionally or not. (That is, so long as the context of that utterance is one in which both of the practices could be felicitously performed.) Since this pair of practices could exist without a given speaker’s knowing about them, the speaker themselves might be in the out-group, i.e., among those competent users of the phrase who are unfamiliar with the narrower usage. On the other hand, the in-group—i.e., those in the audience who *are* familiar

with the narrow usage—can identify the speaker’s utterance as instantiating the narrow practice, even if the speaker had not intended to instantiate it. Thus, the fact that covertly coded speech relies on a pair of practices structured in this way thus explains how speakers might *unintentionally* produce covertly coded speech.

3.1 The Structure of Covertly Coded Practices

I’ve proposed that what enables covertly coded speech is a pair of linguistic practices with a certain structure: both are practices of using the same expression, though they send different messages; both can be performed in (some of) the same contexts; and while one of these practices is familiar to all competent users of the expression, the other is familiar to only some. To see why covertly coded speech depends on a pair of practices structured in this way, let’s revisit the distinctive effects produced by such speech.

First, we noted that covertly coded speech sends (at least) two distinct messages, so there must be (at least) *two* distinct practices in the speaker’s and audience’s shared linguistic community, one practice that sends the ordinary message and another that sends the coded message. Since the ordinary message is sent to all hearers (both in- and out-groups), the practice involving the ordinary message must be a *widespread* one, shared by all competent users of the expression. The coded message, on the other hand, must be sent by a *narrow* practice, one that exists only among some segment of the linguistic community. Next, recall that *covertly* coded speech, as contrasted with overtly coded speech, hides its own codedness by sending, in addition to the coded message, an ordinary one *which seems appropriate* in the discourse context. So the two practices must resemble one another in two ways: they need to not only involve the same expression, but also be deployable in at least some of the same contexts.

In contrast, the practices underlying *overtly* coded speech are not constrained in this way; they may involve different expressions than any widespread practice or be performable only in contexts specific to those narrow practices. For example, the practices of using Pig Latin or technical terms like “supererogatory” and “allocution” are also narrow practices familiar to only some competent English speakers. However, unlike covert codes like “law and order,” such expressions do not also figure in another, widespread practice shared by all competent speakers.

Moreover, not just any pair of widespread and narrow practices will enable covertly coded speech; the two practices must also be structured so that they’re both performable in (at least some of) the same contexts. For some expressions, we might find that there is both a widespread and a narrow practice, but they’re not both deployable in the same context. For example, there’s both an ordinary, widespread use of “title” and a narrow, legal one (as in the certificate of title for a car). Nevertheless, that’s not yet enough to make covertly coded speech possible. Few utterances of “title” are likely to be covertly coded, simply due to the distinct contexts in which the ordinary English usage and the legal usage of “title” can be deployed.²⁵

This structure of practices also explains why speakers can plausibly deny having sent the coded message. The deployability of the widespread practice in the same context, and hence the availability of the ordinary message, is crucial. So long as an utterance is recognizable (in the given context) as *either* of the two practices, a speaker will be able to deny instantiating one of them. Plausible deniability is a feature of covertly coded speech generally, and not just dog

²⁵ A similar constraint holds of the contexts in which covertly coded speech is possible: if only the widespread practice is deployable in a given context, no covertly coded speech can be produced in that context. For example, suppose a speaker attempts to send a covertly coded message by uttering “states’ rights” to an audience composed of scholars of Constitutional law engaged in a conversation about the application of the 10th Amendment to state facemask mandates. In such a context, the covertly coded, racist usage of “states’ rights” might simply not be deployable. I discuss further cases of intended but unrealized covertly coded speech in section 3.2 below.

whistles. Both dog whistles like “law and order” and “inner city” and other covert codes like “Paging Dr. Strong” grant speakers plausible deniability precisely because such expressions are each used in *two* practices.²⁶ Hence plausible deniability will not come with *overtly* coded speech, which involves expressions for which there is no contextually appropriate widespread practice. So, for example, a speaker who utters some expression of Pig Latin would not be able to deny that they’re speaking in code or disown the coded message sent by such an utterance, precisely because there’s no other practice involving that expression, and hence no other message that expression could send.²⁷

3.2 Applying the Practice-Focused Approach

This, then, is the structure of linguistic practices we’ll need to look for, in order to identify covertly coded speech in the world: two practices of using a certain expression, one widespread and one narrow, both deployable in at least some contexts, but sending different messages. When such a pair of practices exists in a community, covertly coded speech (with that expression) is possible. In the next section, I’ll say more about the notion of *practice* I’ve been using so far. But first, I’d like to note the explanatory purchase this account provides.

²⁶ How plausible must this denial be? In some cases, evidence about the speaker’s beliefs and attitudes might make it impossible to disown the coded message. In other cases, particularly low-stakes ones, the mere use of code seems enough to grant the speaker plausible deniability. For example, in the TV show *Penn & Teller: Fool Us*, the magicians Penn Jillette and Teller try to figure out how guests on the show perform their tricks, and then to communicate to the performer how they think their trick was done. To avoid revealing another magician’s secrets, they do so by speaking in code. Despite how easily anyone could “decode” these messages by searching online, there seems to be no backlash against Penn and Teller, suggesting a broad consensus that the codes they use *do* thus allow them to plausibly deny having revealed any secrets.

²⁷ However, this needn’t be the only mechanism responsible for plausible deniability. See Camp (2018), who discusses the covertly coded use of “Dred Scott” as well as other kinds of speech which grant speakers plausible deniability.

On my proposal, the reason Trump and others today can use “law and order” as a covert code is that prior speakers created and maintained the pair of practices involving “law and order”—one practice that’s widely shared by all competent English speakers, and another, racist one practiced by only some. So, to identify the two messages (coded and ordinary) sent by the phrase, we’ll need to identify the two practices (narrow and widespread) in which it figures. That is, the particular coded message sent by Trump’s use of “law and order” in a discourse context concerning anti-police-brutality protests only becomes fully clear once we identify the pre-existing racist practice it instantiates. That practice is one made prominent by Richard Nixon during his 1968 presidential campaign.²⁸ Nixon’s calling for “law and order” in response to Civil Rights protests served not only to dodge the substantive issues at the heart of the protests, but also to disparage the protestors by re-positioning them as mere criminals.²⁹ Moreover, Nixon’s usage was itself not new.³⁰ Indeed, in his usage, “law and order” served many of the same functions and relied on the same stereotypes as it had decades earlier, when the phrase had been used to uphold white supremacy in the post–Civil War Reconstruction era (1865–1877) and to justify the appalling system of convict leasing in the Jim Crow era after that (Haney López, 2014, pp. 38-41).³¹ With the history of this practice in view, we can see that it’s not merely a

²⁸ That Trump’s usage adopts Nixon’s is further evidenced by Trump’s coupling “law and order” with references to the “silent majority,” a phrase also used by Nixon in 1968. Roberts and Jacobs, “Donald Trump Proclaims Himself ‘Law and Order’ Candidate at Republican Convention.”

²⁹ As the legal scholar Ian Haney López (2014, pp. 23-24) explains, Nixon’s use of the phrase “shifted the issue from a defense of white supremacy to a more neutral-seeming concern with ‘order,’ while simultaneously stripping the activists of moral stature.”

³⁰ The legal scholar Michelle Alexander (2012, pp. 40-41) notes that appealing to “law and order” was a widespread strategy of those opposing the Civil Rights movement well before Nixon’s 1968 campaign, starting in the mid-1950s.

³¹ Convict leasing was the practice of state penitentiaries “leasing out” prisoners to private businesses to do heavy, dangerous labor, for which they received no pay. The vast majority of prisoners were Black and had been convicted for petty or specially constructed crimes (such as vagrancy), often on false charges. Convict leasing was explicitly

coincidence that Trump’s usage and the message it sent resemble Nixon’s. Rather, the way that other speakers before Trump have used “law and order” is precisely what made it possible for the phrase to send the particular racist message it did.

A structurally similar account can be given for the example of the 1988 “Weekend Passes” ad (discussed above), which—in the guise of a policy critique—presented viewers with a mugshot of a Black man, Willie Horton, paired with a grisly description of the violent crimes he had allegedly committed. What enabled the ad to function as covertly coded speech was a pair of practices structured in the way I’ve described: one a neutral, ordinary, widespread practice of illustrating a general claim with a single concrete case that serves as a particularly stark example, and the other a racist, stereotyping practice of presenting Black men as innately criminal and as posing a sexual threat to white women. Audiences’ familiarity with the racist practice enabled the ad to prime certain racist attitudes or stereotypes and stoke viewers’ fears. But at the same time, the ad also (at least seemingly) instantiated the non-racial practice of giving a stark example. Thus, insofar as it’s possible for something to instantiate the neutral practice without instantiating the racist practice, the ad’s exploiting both practices created plausible deniability (at least, more plausible deniability than would be available to a speaker who explicitly endorsed the same racist stereotypes).

In order to identify the message sent by Trump’s utterance of “law and order,” or by the display of a Black man’s mugshot in the “Weekend Passes” ad, we asked, not about any and all effects produced by an utterance, but instead about the effects systematically produced by the practice that utterance instantiated. The *message* sent by an utterance, then, is the contribution

permitted by the Constitutional amendment that had abolished slavery, and it existed in some form through the 1940’s. Blackmon, *Slavery by Another Name*.

the utterance makes to a discourse in a certain context, when the utterance is understood as a performance of a certain practice. This contribution encompasses the discursive effects that recur systematically across performances of the same practice in the same context-type—that is, across utterances of the same expression playing the same role in the same context-type. As we’ve seen in the case of “law and order” and the ad “Weekend Passes,” the discursive effects central to a certain practice include not only truth-conditional contributions to the discourse context, but also, for example, affective or evaluative frames. In sum, the message that could be sent by uttering an expression in a given context is a *structural* feature of the practices in which that expression figures; it’s what unites a number of performances, *as* performances of the same practice.

By identifying the conditions in which covertly coded speech is possible, the practice-focused approach also equips us to explain why covertly coded speech is not possible in certain conditions. To illustrate, let’s consider a few cases in which a speaker intends and attempts to produce covertly coded speech but nevertheless fails:³²

- (1) A speaker intends to send a racist covertly coded message by uttering “welfare,” but as it turns out, everyone in the audience is familiar with the word’s narrow racist usage, in addition to its broader, non-racial usage.
- (2) A speaker intends to send a racist covertly coded message by uttering “states’ rights,” but it turns out that no one in their audience is familiar with the narrow, racist usage of this expression, because their sole language is British English.

³² Though she doesn’t take a practice-focused approach, Camp (2018) also focuses on cases where communicated content can’t be adequately captured or explained solely in terms of the speaker’s communicative intention. More careful attention to such phenomena, Camp argues, productively complicates some of our theoretical tools, such as the notions of *common ground* or *what’s said*.

In these two cases, a speaker's utterance does not produce the hidden differential effects distinctive of covertly coded speech. Rather than splitting the audience into in- and out-groups, these utterances send the same messages to all hearers; everyone receives both the coded and ordinary messages in (1), and everyone receives only the ordinary message in (2). Even *intentionally* speaking in covert code is thus only possible before certain kinds of audiences, ones that include at least some members ignorant of and some members familiar with the narrow practice. Insofar as a speaker might misjudge their audience's familiarity with the narrow practice, their communicative intentions alone are not enough to make the utterance a covertly coded one. Recognizing covertly coded speech thus requires careful attention to the social and linguistic structure of an utterance's audience. Whether a certain utterance is covertly coded speech partly depends on something independent of the speaker's communicative intentions—on hearers' differential membership in certain linguistic sub-communities or competence with certain linguistic practices.

Moreover, we'll also need to attend to the social and linguistic structure, not just of the *audience* of an utterance, but of the speaker's and hearers' shared linguistic *community*.³³ For example, consider the following two cases:

(3) A speaker has previously been informed that “welfare” sends a racist covertly coded message, but, forgetting this, intends and attempts to send such a message by uttering “Social Security” instead.

³³ For the purposes of this paper, I'll be working with a somewhat simplified notion of a linguistic community as a set of speakers who can and do understand one another because they've acquired the same patterns of linguistic behavior via socialization. However, our actual linguistic communities are articulated in complex ways, with many, often overlapping sub-communities, fuzzy boundaries, and shifting (and even contested) memberships.

(4) A speaker intends to send a racist covertly coded message by using “sons of bitches” to refer to Black men who are demonstrating against police brutality suffered by Black Americans.³⁴

In these two cases, the differential effects distinctive of covertly coded speech are not produced, because the speaker’s and hearers’ broader linguistic community does not engage in the pair of practices that would make covertly coded speech possible. In (3), the speaker fails to produce covertly coded speech because Americans broadly support Social Security, and there’s no narrow practice among them, in which “Social Security” sends a racist coded message.³⁵ As for (4), a slur like “sons of bitches” with an unambiguous and unambiguously racist referent doesn’t send any ordinary message in addition to its racist, derogatory one, and thus doesn’t produce the differential effects distinctive of covertly coded speech. So, despite the speakers’ communicative intentions, everyone in the audience receives the same message—a neutral or positive one in scenario (3) and a racist, derogatory one in (4). As in cases (1) and (2), we see that speaking in covert code, even intentionally, at least partly depends on something independent of the speaker’s communicative intentions—in these cases, on facts about their linguistic community as a whole. Taken together, these four cases show that differential effects, hiddenness, and

³⁴ When Donald Trump used such language in reference to NFL players in September 2017, a number of observers, including Hillary Clinton, called it a racist dog whistle. For example, Hillary Clinton called it a “dog whistle to his base.” Ian Schwartz, “Hillary Clinton: Quite Telling That Trump Is Willing To Attack Black Athletes, ‘Dog Whistle’ To Base,” RealClearPolitics, Sept. 25, 2017, https://www.realclearpolitics.com/video/2017/09/25/hillary_clinton_quite_telling_that_trump_is_willing_to_attack_black_athletes_dog_whistle_to_base.html.

³⁵ Winter (2006) summarizes the empirical data on differences between Americans’ support for welfare and for Social Security, and argues that this difference is attributable to how the two have been racialized—while “welfare” is associated with blackness (and hence with racist stereotypes about Black Americans), Social Security is associated with whiteness.

recurrence—the three key features of covertly coded speech—arise only if the audience, as well as the speaker’s and audience’s shared linguistic community, are structured in the right ways.

4. Focusing on Practices

The practice-focused approach I’ve been advocating builds on the uncontroversial observation that much communication relies, not just on properties of individual speakers, but on interlocutors’ shared membership in linguistic communities and their shared practices. Any one speaker is embedded in a community of other speakers, and any one utterance in a practice performed by others. Such dependence on other speakers in our linguistic communities is not unique to coded speech. Consider, for example, how difficult it is to successfully use “begs the question” in its “correct” sense outside an academic context; one will simply not be understood as saying something about the logical structure of an argument. The uncontroversial observation is that speakers’ ability to use words in certain ways is constrained by how those words are commonly used by others in their linguistic community. But taking this observation seriously means that our theories of communicated content must be sensitive to the structure of practices in a speaker’s linguistic community—and not just to facts about individual speakers—whether we’re theorizing covertly coded speech or communicated content more generally.

I’m certainly not the first to note the inadequacy of thinking of speakers in isolation. Indeed, Jason Stanley’s (2015) account of dog-whistle speech emphasizes the role of a speaker’s community—and especially of past speakers—in associating certain expressions with certain

social meanings.³⁶ Stanley argues that this associative process is the very same mechanism by which expressions acquire their ordinary conventional meaning (p. 138). For example, this same mechanism is responsible both for the conventional meaning of an ordinary expression like “dog” and for the racist, coded message sent by a dog whistle like “welfare” (ibid.). So, on Stanley’s view, the racist message sent by a phrase like “law and order” is a negative social meaning encoded by the expression and thereby communicated by all subsequent utterances, with or without the speaker’s intent.³⁷

While I share Stanley’s emphasis on the role of past speakers in associating certain expressions with certain messages, the associative mechanism alone can’t account for all of the distinctive features of covertly coded speech. In particular, neither the mechanism of association nor the notion of social meaning suffice to explain how an utterance can be, not just coded, but covertly coded. As we’ve seen in the scenarios discussed in the previous section, both a speaker’s audience and their linguistic community need to be structured in certain ways, in order for that speaker to produce not just differential effects, but hidden differential effects which recur systematically across similar utterances. Identifying the conditions that enable covertly coded speech requires identifying both the history of usage and the resulting *structure* of covertly coded practices. While Stanley’s account rightly assigns the history of usage a central role, my proposal

³⁶ Roughly, the *social meaning* encompasses the beliefs, expectations, norms, and ideals that members of a community typically associate with an institution, practice, or linguistic expression, as the product of their shared socialization into those institutions, practices, or into the use of that language. For a more detailed discussion of social meaning, see Stanley (2015), pp. 157-62, 167-68, and Haslanger (2018).

³⁷ Though Stanley argues that, through association, the negative social meaning becomes part of the *content* of an expression like “law and order,” he’s careful to note that such contents needn’t always be propositional (p. 146). They could also take the form of, for instance, emotional or normative effects, which can be better modeled by an expressivist semantics, as contributions to a conversation which introduce *preference orderings* that rank possible worlds (rather than as propositions that rule certain possible worlds in or out).

fills in the crucial details by identifying the particular structure of practices that enables covertly coded speech.

In her recent work on dog whistles, Jennifer Saul (2018) also explicitly recognizes how a linguistic community's pre-existing practices might shape what communicative options are available to a given speaker.³⁸ Having identified the distinctive effects produced by dog-whistle expressions like “inner city,” Saul turns to consider how a speaker could *unintentionally* produce such effects. She addresses this question mainly as it pertains to dog whistles that raise certain attitudes to salience without the awareness of hearers thus affected—i.e., a narrower category than covertly coded speech.³⁹ However, the two responses Saul considers generalize, and her discussion of the question is thus illuminating as regards to covertly coded speech more generally.

How, then, could a speaker *unintentionally* produce such effects? Saul considers two answers to this question (pp. 377-78). The first, which she ultimately endorses, replies that there's nothing unique to dog-whistle speech here. Speakers can unintentionally produce the same dog-whistling effects that other speakers produce intentionally, much as speakers can unintentionally insult someone by uttering some words which another speaker uses to

³⁸ See also Saul and Drainville (forthcoming, §1.3), noting that, when there's an established usage, unintentional dog-whistle speech can take either of the two forms identified by Saul, overt (i.e., expressions that were designed as codes and that function by signaling a particular content of which the in-group is aware) and covert (i.e., expressions that function by raising certain attitudes to salience, without the awareness of those affected).

³⁹ Saul calls this type of speech a “covert dogwhistle,” though her usage of “covert” differs from mine. Saul's covert dogwhistles would be a proper subset of the category of covertly coded speech identified in section 2 above, because her distinction between “overt” and “covert” tracks only the affected hearers' conscious awareness of what I call the coded message. However, I don't distinguish between coded messages that are received consciously or unconsciously in this paper. As I use the term, “covertly coded” speech is a subset of coded speech distinguished by the three features identified in section 2—differential effects, hiddenness, and recurrence.

intentionally insult someone. Unintentional dog-whistle speech is thus just speech that produces similarly “pernicious unintended perlocutionary effects” (p. 377).

The second option considered by Saul identifies some other past speaker’s communicative intention as the one that makes something a dog whistle. On this option, speakers may unintentionally give voice to a dog whistle, but the responsibility for the pernicious effects ultimately lies with some other, past speaker. This option, Saul points out, emphasizes the manipulative aspect of dog whistle speech, by locating the manipulation that makes unintentional dog whistling possible: “Those who create the initial covert dogwhistles are very good at attaching pernicious associations to words and images (and possibly other things as well) and sending them out into the world in the hope that they will be taken up and used by others, bringing with them these associations” (p. 378).

Saul rejects this second option because it would seem to commit one to saying that the “real” speaker is not the person producing some sounds or text, but the ones who dog whistled intentionally at some previous time, the ones whose intentions attached the pernicious associations to the words (p. 378). This option thus not only involves a counter-intuitive notion of “speaker,” but also misplaces agency—and thus, responsibility—away from the producer of the utterance. And those giving voice to a dog whistle “really are the speakers,” Saul writes, “and they need to be thought of as such, and held accountable for the effects of their speech” (p. 378).

Though I share Saul's concern that speakers aren't let off the hook too easily for harmful, derogatory speech, I don't think this is this is an unavoidable consequence of the second option. In fact, the practice-focused approach I've been advocating points to a way of developing that second option in a way that avoids the problems identified by Saul. The focus on linguistic practices draws our attention to the ways in which (and the reasons why) the message sent by an

utterance is not solely determined by the speaker. But this doesn't foreclose the possibility of holding speakers responsible; it highlights alternatives for what we might hold speakers responsible for, besides just meaning this or that thing by their words. For example, from the practice-focused approach, we can explain why and how we might hold speakers responsible for maintaining or extending a practice which enables speech we find reprehensible. Or, for another example, the practice-focused approach helps explain why speakers positioned in certain way (as politicians are) should be responsible for a certain amount of due diligence, when it comes to considering their audiences in choosing their words. This approach can thus help to connect speakers' responsibilities as speakers to their responsibilities as, say, public representatives: precisely because coded messages recur systematically, we could explain why the greater concern with covertly coded speech lies not in any particular speaker's having a reprehensible intention, but rather in the very existence of certain practices which allow public officials to violate norms of transparency and even-handedness in their treatment of their audiences.

Of course, I cannot offer a full account of speaker responsibility here. Rather, I raise these possibilities in order to re-open something like the second option considered by Saul, where the use to which past speakers put a word shapes the uses it can serve in the future. That option is one which, despite Saul's reservations, can be productively developed in terms of the practice-focused approach. In the previous section, we saw how this approach helps to identify the conditions in which covertly coded speech is and isn't possible. In the rest of this section, I detail the practice-focused approach more thoroughly. To show how we can fruitfully deploy this strategy and identify practices, their history, and their structure, I'll need to say a bit more about what it means to view an utterance as a performance of a practice, and more generally about social practices themselves.

4.1 Practice-Focused Approaches to Communication

The approach I've been advocating prioritizes giving explanations in terms of *communal practices* of communication, rather than in terms of individual mental states (such as communicative intentions, as for example in Grice, 1957) or abstract formal entities (such as interpretation functions or sets of possible worlds, as for example in Lewis, 1969). This approach thus emphasizes the social, political, and historical aspects of linguistic practices. Thinking of utterances as performances of certain practices enables us to identify how sociopolitical historical contexts shape what's said or what speech act is performed. I'm not alone in arguing that a practice-focused approach is essential to understanding the mechanisms underlying certain kinds of speech. For instance, Lynne Tirrell advocates just such an approach in her work, and has applied it to genocidal speech (2012) and slurs (1999), to name but two examples. In her study of the kind of dehumanizing speech that leads up to a genocide, Tirrell argues that, to accurately identify the function of such language, we must look more carefully at the sociopolitical practices in which it's embedded. Further, when we take that practice-focused approach, we see how that function can be achieved even without the speaker's intending it (Tirrell, 2012, pp. 187-88).

The linguist Sally McConnell-Ginet has also advocated and refined a practice-focused approach, applying it to investigate the dynamics of language and power in connection with, for example, gender and sexuality (2011) and social group labels (2020). Examining linguistic phenomena through the lens of *communities of practice* has been a particularly influential aspect of this work. Adopting the notion from the anthropologists Jean Lave and Etienne Wenger (1991), Penelope Eckert and Sally McConnell-Ginet introduced the notion into sociolinguistic technique in their work on gendered language use (e.g., differences between men's and women's

speech) and the language of gender (i.e., terms like “man” and “woman”, “Mr.” and “Mrs.”).

Eckert and McConnell-Ginet (1992) argue that these linguistic practices are not fully intelligible independently of the other social practices with which they’re interwoven. As such, they’re best understood as communities of practice, whose membership is constituted, not by inherent or assigned properties of individuals, but rather by their shared projects, techniques, and practices—including their shared linguistic practices.

This notion of communities of practice has proven useful in philosophy of language as well. For example, Luvell Anderson (2018) applies it in his recent work on reappropriated slurs, in order to explain why some reappropriated slurs (such as the N-word) can only be used by speakers who are members of the group targeted by that slur. Anderson argues that, like other speech acts, the N-word can only be used in its reappropriated sense by speakers who are members of a particular community of practice.⁴⁰

Unsurprisingly, the utility of a practice-focused approach is particularly clear when we focus on public or political speech. So, for instance, in his account of propagandistic speech, Jason Stanley (2015) repeatedly emphasizes the historically, politically embedded linguistic practices that enable and constrain what we mean or do with words. Stanley points out that our speech practices encode not only linguistic meaning, but social meaning as well. He further argues that theories of communication shouldn’t ignore social meaning, for the social meanings

⁴⁰ See also Hess (2020), who offers another practice-focused account of reappropriated slurs, employing a somewhat different conception of practice, in order to explain some of the same phenomena as Anderson.

of our words—and not only their truth-conditional contributions—impact hearers’ mental and affective states and their behavior.⁴¹

More recently, David Beaver and Jason Stanley (forthcoming) amplify this focus on politically significant language, and offer a critique of and an alternative to theories which conceive of communication as the exchange of contents. Beaver and Stanley’s model of communication prioritizes communicative *practices*, rather than instances of communication considered in isolation, and focuses on communicative *effects* rather than contents. This shift in priorities, they argue, allows them to identify the mechanisms responsible for the (often pernicious) effects on audiences’ emotions, reasoning, and behavior produced by what Beaver and Stanley call *hustle*—non-cooperative, manipulative speech such as propaganda or dog whistling. Beaver and Stanley model these mechanisms in terms of the *resonances* carried by words and acts, resonances to which audiences are attuned as a result of socialization—that is, resonances that are present (or not) to certain hearers because they’re members of certain communities. The resulting theory is one which allows Beaver and Stanley to explain, not just how one speaker can communicate something truth-apt to one hearer, but also how affective and behavioral alignments, as well as ideologies more generally, spread in mass audiences and through dispersed communities across timeframes ranging from minutes to decades.

Besides foregrounding the sociopolitical and historical aspects of discursive context, the practice-focused approach emphasizes the ways in which our shared practices (and not just our private mental states) imbue our actions with publicly accessible meaning. Something can only be a vote, baptism, marriage, or testimony if there exists the right kind of practice; otherwise

⁴¹ For example, social meanings can affect reasoning and behavior by making a derogatory stereotype salient, prompting certain emotional responses, or making certain norms, values, or frames salient. For a more detailed survey of the range of propagandistic mechanisms, see Stanley (2015), 139-72, and Quaranto and Stanley (2021).

these would be mere bodily movements and sounds.⁴² Linguistic practices similarly imbue actions with meanings, making it possible for an utterance to send a certain message. This has been the central tenet of various approaches to foundational questions about meaning, representation, and intentionality. The American pragmatists, the later Wittgenstein, and externalists like Putnam and Burge are just a few examples. John Dewey (1925, p. 148), for instance, describes meaning as a type of social or joint action, writing that “The heart of language is not ‘expression’ of something antecedent, much less expression of antecedent thought. It is communication; the establishment of cooperation in an activity in which there are partners.”⁴³ A similar tenet is expressed in Wittgenstein’s (1953, §§198, 199) influential remarks describing language use as involving, not just mastery of a technique, but also participation in an institution or a practice. And although they aren’t always cast explicitly in practice terms, externalist accounts such as Putnam (1975) and Burge (1979) also explain meaning in terms of shared, public practices (rather than private mental states).

But perhaps the most committed and influential practice-focused approach to foundational and applied questions in philosophy of language is the inferentialist project developed in Brandom (1994), and refined and extended in works such as Lance and O’Leary-Hawthorne (1997), Tirrell (1999) and (2012), Kukla and Lance (2009), and Kukla (2014), to name only a few. Brandom (1994, pp. 26–42) argues that, if practices imbue acts with meaning, then they can’t be mere regularities in our behavior or dispositions. He thus advances a

⁴² This is of course not a new point, nor is it confined to the practice-focused approaches I’ve described above. For instance, this point is also central in Austin (1962) and in Searle (1995).

⁴³ See also Dewey (1925), chapter 5 more generally for his pragmatist approach to language and communication. G.H. Mead’s (1934) account of thought and intentionality also emphasizes communal representational practices. Mead (pp. 42-82) is particularly sensitive to the interactive aspect of communication, and the ways in which this interactivity depends on agents’ sharing linguistic dispositions, which in turn is possible only because these dispositions have been acquired in a social context.

normative conception of discursive practice, one which grounds the notions of meaning, reference, and truth in his pragmatic, inferentialist theory. Drawing on Sellars, Brandom conceives of discursive practice as a “space of reasons” where we can ask for and offer justification (p. 89). On this view, asserting is a move which commits the speaker to justify their claim, and thereby licenses hearers to make certain inferences (pp. 172-75).⁴⁴

Though it’s not the line I’ve developed in this paper, covertly coded speech could be analyzed in these inferentialist terms, saying for instance that the in-group attributes different commitments to the speaker than the out-group does. While such an account is compatible with my proposal in this paper, I have not pursued it here because my goal has been primarily to identify the *enabling conditions* of covertly coded speech. While the emphasis on normative inferential relations allows Brandom to account for fundamental notions like reference and objective truth, inferential relations alone don’t fully explain all the functions, features, or effects of a given type of speech. Rather, an inferentialist account must also attend to the sociopolitical and historical contexts in which speech is produced. Of course, this point is familiar to inferentialists; indeed, as we’ve already seen, Lynne Tirrell argues for just this point, in presenting her inferentialist account of genocidal speech. But the covertness and deniability of covertly coded speech pose particular explanatory challenges, and highlight the importance of

⁴⁴ That is, in keeping conversational score, speakers and hearers assign normative statuses (commitments and entitlements) to themselves and to one another, and they do so against a background of normative discursive practices. Brandom (p. 159) argues that these practices are not merely regularities in how we in fact keep conversational score; rather, they prescribe how we *should* attribute commitments to speakers and entitlements to hearers.

thinking in terms of practices as wholes.⁴⁵ That is, to explain what makes covertly coded speech possible, or to identify the conditions in which covertly coded speech produces its distinctive effects, we must identify the particular *structure* of covertly coded practices. Only that structure—the pair of practices identified in the previous section—allows us to explain the mechanism behind the recurring, hidden differential effects produced by covertly coded speech.

4.2 Identifying Practices

We've seen that a practice-focused approach involves identifying the history and structure of linguistic practices. This in turn requires knowing what it is we're looking for, when we look for a linguistic practice. However, identifying linguistic practices *doesn't* require knowing the necessary and sufficient conditions for something's being a linguistic practice. Though theorists employing a practice-focused approach disagree about the definition or features of social practices, the notion is not solely a technical one. Rather, as agents in many social practices ourselves, we have a substantive pre-theoretical notion of social practice, one which routinely enables us to identify practices and the conditions for their performance.

As a widely shared starting point, we can say that a linguistic practice is a socially instituted pattern or network of the linguistic behaviors and dispositions of the members of a linguistic community.⁴⁶ Of course, there are many different ways to go from here, and the

⁴⁵ So, while my account of covertly coded speech is ultimately compatible with an inferentialist account, there's no lossless translation of my proposal into claims solely about individuals' ascriptions of discursive commitments. Rather, my proposal could complement such an account, by saying why and when different audience members ascribe commitments differently, and by explaining how the differences in their scorekeeping arise from the structure of practices in which audience members are differentially positioned.

⁴⁶ In this paper I focus on speakers' and hearers' behavior, rather their mental representations, propositional attitudes, or affective states. Of course, the particular actions that make up a practice typically involve and may sometimes require certain representational mental states (beliefs, intentions, and expectations) and, especially for

literature on social practices offers a number of competing conceptions of what our practices are like, what it means to instantiate or participate in them, and what's required for competence or membership in a practice.⁴⁷ Theorists disagree, for instance, whether practices (and other collective phenomena) reduce to individual agents' actions or mental states. Brandom (1994, pp. 38–39) argues that discursive practices must ultimately supervene on individual attitudes, because only individuals, and not groups, can be coherently described as holding normative attitudes.⁴⁸ Others, like Bourdieu (1977) and Haslanger (2016), argue against such a reduction. Haslanger (2016) argues that attempts to reduce the collective to the individual overlook the ways in which individuals' attitudes and actions in fact often depend on the social, insofar as our communities' practices make certain attitudes and actions possible (or impossible).

There's also disagreement about the relations between social practices and other collective phenomena, such as conventions or social norms. So while some, like Tuomela (2002), parse practices in terms of conventions in the sense of Lewis (1969), others, including Gilbert (1989) and Haslanger (2016), have argued that Lewisian conventions don't accurately capture some of the key features of social practices.⁴⁹ Manipulative, non-transparent practices seem to support the latter view, given that Lewisian conventions involve coordination among

derogatory dog whistles, performances of a practice and their reception are often accompanied by certain affective states. A more nuanced account of what our practices are like should account for the relations between performances of a practice and practitioners' beliefs, attitudes, and emotions. However, providing such an account is out of the scope of this paper, as my goal here is to identify the structure of practices that enables covertly coded speech. And as the worked-out examples in this and the previous section show, identifying covertly coded practices doesn't require full information about speakers' and hearers' mental states. We can identify the relevant structure of practices on the basis of empirical and historical data about the effects of utterances.

⁴⁷ To name but a few influential accounts of social practice: Bourdieu (1977), Giddens (1984), Gilbert (1989), Schatzki (1996), Tuomela (2002), Hanna & Harrison (2004), Rouse (2015), Haslanger (2016) and (2018), and McMillan (2018).

⁴⁸ Searle (1995) also takes collective phenomena to ultimately reduce to individual mental states.

⁴⁹ Haslanger (2016, p. 126n9) argues that, unlike Lewis' conventions, practices aren't always arbitrary, needn't be common knowledge among participants, and aren't always mutually advantageous.

equally positioned agents. Moreover, as we've seen, attention to historical, sociopolitical context is a central feature of the practice-focused approach. It's unclear whether such attention is ultimately compatible with the abstractions of Lewis's definition of conventions. For instance, per Lewis (1969), a necessary condition is that the actual regularity in a community's behavior is not the only possible one. For example, in developing our driving practices, we could have chosen to drive on the left side of the road instead of the right. But whether a given expression can be used to send covertly coded messages is not so arbitrary; the structure of practices and the history of usage matter for determining what message(s) the expression can send.

Given the number and complexity of these disagreements, I can't hope to resolve them here. Instead, I rely on a more minimal, pre-theoretic notion, which we'll see is nevertheless sufficiently rich for our purposes here. As social creatures and practitioners, we're already equipped with a pre-theoretic notion of social practice. We can typically say which practices we do and don't participate in, even if we can't state the necessary and sufficient conditions for that participation. For instance, we have no trouble recognizing that contemporary English speakers don't participate in politeness practices marked by the T-V distinction for second-person pronouns, while speakers of contemporary French, German, and Russian do. Our ordinary notion of practice also recognizes multiple levels of generality. In speaking of the T-V distinction, we refer to the more general practice shared by French, German, and Russian speakers. Alternatively, we can refer to the more specific practices of each of these linguistic communities, by speaking of the *tu/vous* or *du/Sie* or *ты/вы* distinction. And we can take an even more fine-

grained approach, distinguishing for example between the *vous* practices of different generations of French speakers.

So while it may be minimal, the pre-theoretic notion of practice is not uninformative. In fact, it's sufficiently rich to distinguish social practices from at least some other collective phenomena. When we think of French speakers' different uses of *tu* and *vous* as their performing different practices, we recognize that the observed patterns of behavior aren't just coincidental. Unlike mere statistical regularities, socially instituted practices exist not merely by chance, nor as a brute natural fact. For example, though most people will flinch if startled, we don't engage in a flinching *practice*.

Our pre-theoretic notion of social practice also incorporates various key properties of social practices. For instance, practices needn't be explicitly instituted or consciously, deliberately performed. Consider American practices of maintaining personal space, such as how we arrange ourselves in an elevator or where we sit on a bus. As participants in these practices, we have no trouble recognizing that we participate in them even though we never explicitly agreed to behave this way, nor were we (typically) explicitly instructed what to do. Rather, when there's a practice, a community's behaviors and dispositions pattern in a certain way because its members learn to behave that way through socialization, by interacting with other members of their community and thereby joining in their community's practices.⁵⁰ Crucially, in our pre-theoretic understanding of social practices, we also recognize that our behavior can exemplify a certain practice, even without our thinking of it as such, and sometimes, as in the personal space

⁵⁰ For a more detailed account of how such socialization might give rise to practices, see Hanna & Harrison (2004), pp. 169–72.

example, even without our thinking about the behavior at all. That is, as competent practitioners, we recognize that we might perform some practices unintentionally.

5. Incorporating the Practice-Focused Approach

Thus far, I've identified the enabling conditions for covertly coded speech in a certain structure of practices—namely, a pair of practices involving the same expression but sending different messages, both deployable in at least some contexts, but one widely shared by all competent users of the expression and the other only by some. I've also argued that taking a practice-focused approach is crucial for a deeper understanding of covertly coded speech. That said, the practice-focused approach is not the only available lens, and depending on our explanatory goals, we might take higher- or lower-level views. For instance, we could consider an utterance as just a physical event, which might be suitable if we're interested in measuring the frequency and pitch of the speaker's vocal productions. But this level would not be appropriate for comparing, for example, the different strategies speakers employ to make their needs or desires known to intimate partners and to strangers. For that inquiry, it would be more appropriate to view the utterance as an intentional act or an expression of inner states. At a higher level, we can think of an utterance, not solely as an individual action, but as a performance of a practice also instantiated by other agents at other times and places. This is the

level I've focused on here, and the level that has enabled me to identify what makes covertly coded speech possible.⁵¹

Not only do different goals call for different approaches, but even one and the same explanatory goal might require multiple approaches taking different perspectives at different levels. The practice-focused approach deployed and defended in this paper is thus not meant to replace all other lenses through which we might view communication. Certainly utterances sometimes convey a message, not by instantiating a particular pre-existing practice, but more idiosyncratically, for example via Gricean conversational implicature. By itself, the practice-lens might not be very productive for understanding such utterances, for its higher-level focus on what's common to all the performances of a practice might obscure some of the one-off particularities crucial for explaining how that one particular message was sent. In contrast, the covert codes which I've focused on here—expressions like “law and order” and “inner city”—send messages that systematically recur across all utterances of the same expression in the same context-type. The higher-level view of the practice-focused approach is particularly well-suited to explaining the mechanisms responsible for this kind of message.

That said, the practice-focused approach to covertly coded speech needn't rule out all other accounts of covertly coded speech, for, as I'll show in this section, it can be productively assimilated even into views of communicated content which prioritize individual mental states. In this section, I'll motivate such an assimilation and indicate how it might look for two kinds of

⁵¹ Of course, a practice-focused approach is not the only alternative to an individualistic view that focuses on speakers' or hearers' mental states. Indeed, I've referred to several such alternatives in the course of this paper, including Saul's (2018) account of unintentional dog-whistling and Camp's (2018) analysis of insinuation. There are also other non-individualist alternatives which I have not been able to address here. To name but a few, covertly coded speech might be analyzed in terms of conventionalized implicatures, conventionalized indirect speech acts (Asher & Lascarides, 2001), speech acts that change the context in which they're performed (Sbisà, 2002), speakers' direct intentions and hearers' imaginative engagement (Lepore & Stone, 2015), or discourse structure (Roberts, 2018).

accounts, one in terms of Gricean communicative intentions, and the second—proposed in Khoo (2017) and further developed in Henderson and McCready (2019)—in terms of hearers’ beliefs and inferences.

5.1 Speech Practices and Gricean Accounts

Despite any apparent incompatibility, even intention-based theories of communicated content, such as Grice (1957) can and should incorporate the kinds of community-level facts I’ve identified. To illustrate how this incorporation might work, I draw on Herbert Clark’s (1992, pp. 255-274) account of *disguisement*, a speech act closely resembling what I call covertly coded speech. Like covertly coded speech, disguised speech not only conceals from some hearers what is said to others, but also masks that concealment by misdirection. The speaker achieves this misdirection by exploiting the differences between what different hearers take to be the common ground (which in Clark’s usage refers to the set of beliefs held in common by the speaker and hearers and known by them to be held in common⁵²).

In Clark’s example (pp. 267-68), it’s common knowledge among the speaker (Alan), their addressee (Barbara), and the overhearer (Oscar) that there is both a Beethoven concert and a lecture by Louis Levesque occurring that evening. But only Alan and Barbara, and not Oscar, know that Alan and Barbara jokingly use “Louis” to refer to Ludwig Beethoven, and that Alan and Barbara plan to attend the Beethoven concert together. Wanting to hide their plans from

⁵² In other words, the common ground between a speaker and hearer is the set of propositions which are *common knowledge* between them, in the sense of common knowledge defined by Lewis (1969, p. 56) and refined by Schiffer (1972), where some proposition *p* is common knowledge between S and H (or, in Schiffer’s terms, *mutual knowledge*) if and only if S and H both know that *p*, both know that the other knows that *p*, and so on ad infinitum. Thus, the notion of common ground Clark uses is also the one employed in Stalnaker (1978), and closely related to, but also more precise than, the earlier notions of common ground appearing in Karttunen (1974) and Stalnaker (1974).

Oscar, but unable to get out of his earshot, Alan asks, “Do you have ten dollars for your Louis ticket?” Based on her knowledge of her and Alan’s plans *and their usage of “Louis,”* Barbara interprets “Louis ticket” as referring to a ticket to the Beethoven concert. An outsider to Alan and Barbara’s use of “Louis,” Oscar is misdirected into a different interpretation of the utterance. On the basis of his common ground with Alan and Barbara, which includes only the beliefs that there is both a Beethoven concert and a Louis Levesque lecture that evening, Oscar interprets “Louis ticket” as referring to the Louis Levesque lecture.

As in covertly coded speech, in disguised speech like Alan’s question about the “Louis ticket,” all hearers receive some message, but not the same message as one another. Lacking some relevant insider knowledge and hence relying on a different common ground to interpret the speaker, some hearers infer a different speaker meaning than others. And, again as in covertly coded speech, the fact that the overhearer is nevertheless able to recover *some* interpretation hides the fact that the utterance is coded. Recovering *some* relevant speaker meaning, the overhearer doesn’t have reason to suspect that the speaker intended other hearers to infer a rather different meaning.⁵³

As Clark describes it, disguisement works via the general Gricean mechanism of hearers’ reasoning about the speaker’s beliefs and intentions. However, this doesn’t mean that an account like Clark’s rules out unintentional performances. Hearers might mistakenly take the speaker to have certain intentions and beliefs, and so infer and attribute to the speaker a disguised meaning which the speaker didn’t actually intend. For example, rather than intending to conceal from Oscar his plans with Barbara, Alan might utter the same question and produce the same

⁵³ However, disguisement, as Clark describes it, does require the speaker to intend two different speaker meanings to be received by the two different hearers; covertly coded speech as I’ve described it doesn’t require or assume this to be the case.

disguised speech, while simply forgetting that Oscar doesn't share their joking usage of "Louis" to refer to Beethoven.

Even so, an account like Clark's won't yet capture all covertly coded speech. As Clark himself notes, disguise is typically only possible among close friends, because it relies on the speaker's and hearers' having detailed knowledge about one another's beliefs. So, unless speakers and hearers could also rely on some other, more publicly available shared background, Clark's account of disguise could not explain the possibility of covertly coded speech in the public sphere, before large, diverse audiences unknown to the speaker.

Clark does suggest a way to incorporate such a publicly available shared background, and thus to explain disguise among strangers. He proposes extending the notion of the common ground to include a *communal common ground* (p. 257). Like common ground more generally, the communal common ground is the set of beliefs that are common or mutual knowledge between speaker and hearer (i.e., a set of propositions they each believe, believe that the other believes, and so on). But unlike their *personal common ground*, which consists of beliefs the speaker and hearer "have inferred from personal experience with each other," the communal common ground consists of beliefs the speaker and hearer "take to be universally held in the communities to which they mutually believe they both belong" (ibid.). The communal common ground is thus formed on the basis of community membership or social identity, such as a shared language, nationality, education, place of residence, occupation, religion, hobby, or subculture (ibid.). To craft disguised utterances, speakers can thus rely on beliefs or experiences shared in virtue of a shared social identity, and not just in virtue of intimate knowledge of one another's beliefs and intentions. A lawyer, for instance, has various beliefs about what lawyers believe, including mutual knowledge of certain narrow linguistic practices such as technical jargon. In

conversation with another lawyer, they can rely on those beliefs in order to design or interpret utterances.

Positing a communal common ground, as Clark does, is precisely the assimilation I referred to above, the assimilation of a practice-focused approach into an intentionalist view. Indeed, this notion of a communal common ground aims to capture the same shared background that I've described as speakers' and hearers' competence with communal and sub-communal linguistic practices. It's no accident that in Clark's example, the speaker (Alan) relies on a very narrow practice—his and Barbara's pre-existing joking usage of "Louis." Whether we describe it in terms of a shared practice or a communal common ground, what enables covertly coded speech (or disguise) is a shared, interpersonal aspect of language.⁵⁴

It's telling that, in developing a Gricean account of disguise, Clark ultimately appeals to facts about communities' and sub-communities' practices, not just about individual speakers' and hearers' mental states. This shows not only that a practice-focused approach is not only compatible with individualistic, mentalistic accounts of communication, but also that practices play an important explanatory role, one that can't be filled solely by facts about individuals taken in isolation.

5.2 Speech Practices and Hearers' Beliefs

In his recent work on coded speech, Justin Khoo (2017) offers a rather different explanation of how and why in-group hearers receive a coded message. Like Clark, Khoo proposes that what I've called coded messages are *inferences* that some hearers draw from the

⁵⁴ So, a reader compelled by an intentionalist view like Clark's could interpret my description of the pair of practices underlying covertly coded speech as identifying the distinctive structural features of the communal common grounds which speakers rely on to produce covertly coded speech.

speaker's utterance. But unlike Clark, Khoo argues that these inferences rely on beliefs that hearers have prior to and independently of the utterance—namely, beliefs about the world or about who typically uses the expressions in question, rather than beliefs about the common ground or the speaker's communicative intentions (pp. 47-48). So, for example, what explains “inner city” sending a racist coded message to some is their having a pre-existing belief like “The inner city is mostly populated by poor African Americans” (p. 47). The in-group then infers a racist message on the basis of this belief. Relatedly, Clark and Khoo locate the operant inferences at different stages in the interpretive process: whereas in Clark's account of disguise, the inferences drawn by hearers are part of the process of interpreting the speaker's words, in Khoo's analysis of coded speech, hearers rely on the utterance's content in order to infer the coded message. That is, the inferences that Khoo identifies occur *after* the process of semantic interpretation is complete. Thus, these inferences are a pragmatic result of hearers' pre-existing beliefs, rather than a conventionalized part of the coded expression's meaning.

Khoo argues that the decisive evidence for this conclusion is that the coded meaning cannot be isolated. That is, even when their ordinary meaning is explicitly controverted, covert codes never send only their coded message. Khoo points out, for example, that “Smith is an *inner city* pastor who is from, works, and lives in the *suburbs*” is infelicitous (p. 45, emphasis added). If “inner city” semantically encoded some racial content, then there would be an available interpretation of “inner city” that entailed claims only about Smith's race, and not about geography. On such an interpretation, Smith's being “an inner city pastor” would be compatible with his Smith's origins, work, and home being in the suburbs, and we wouldn't find the sentence infelicitous. Thus, for Khoo, the infelicity we in fact observe shows that, whatever

effects the coded, racist message sent by “inner city” has on the discourse, the racist message is not semantically encoded content.

Khoo’s account can rather straightforwardly encompass unintentional dog-whistling is possible. After all, coded messages are the product of hearers’ inferences based on pre-existing beliefs, and such inferences can be triggered without the speaker’s intending to send a coded message (pp. 48, 49). However, Khoo’s account also overgeneralizes the conditions in which such inferences should occur. As Khoo himself notes, his account predicts that uttering any expression that’s co-extensional with “inner city” would produce the same effect as uttering “inner city,” triggering the same inferences from the same pre-existing racist beliefs (p. 50). But this prediction is not borne out. For example, the extent to which Americans express support for certain federal aid programs varies with question wording. An analysis of multiple surveys found that “on average support for more *assistance for the poor* is 39 percentage points higher than for *welfare*” (Smith, 1987, p. 76, emphasis added). Another survey found that white Americans support government spending on “Social Security” vastly more than they do “welfare” (Winter, 2006, p. 406).⁵⁵ Moreover, the lack of support for “welfare” appears despite the fact that 68% of Americans agree that “government must see that no one is without food, clothing, or shelter” and 78% endorse “their own tax dollars being used to help pay for food stamps and other assistance for the poor” (cited in Haney López, 2014, p. 60). The empirical data suggests that “welfare” elicits responses which seemingly co-extensional phrases like “assistance to the poor” do not. So, if “welfare” produces its differential effects by triggering inferences from pre-existing beliefs,

⁵⁵ More specifically, Winter (p. 406) reports that “On a scale from zero to one, support for Social Security spending averages 0.745..., or just about exactly midway between the ‘increase’ and ‘keep it the same’ responses. ... In contrast—and not surprisingly—support for spending on welfare and on food stamps is much lower among whites (average of 0.31 and 0.38, respectively).”

then those beliefs must be (at least partly) about the *expression* “welfare” itself, for which we don’t have good evidence.

Khoo proposes that such data suggest extending the account to include *framing effects*, in order to explain why certain inferences are triggered by, for example, “welfare” but not by “assistance to the poor.” That is, if we take on board the principle that “A rational individual may believe *x* is *F* when thinking of *x* one way, and believe that *x* is not *F* when thinking of *x* in a different way” (p. 51), then the observed difference between co-extensional phrases like “food stamps and assistance for the poor” and “welfare” is not a counter-example to Khoo’s account. While the appeal to framing effects may deflect such counter-examples, it doesn’t explain why there are such differences between inferences triggered by co-extensional phrases like “welfare” and “assistance for the poor.” Nor does Khoo’s account aim to explain these difference, for on his view, they’re a matter of psychology and not language.

Pace Khoo, I think that an account of coded speech should be able to explain, not only the difference in effects between “welfare” and “assistance for the poor,” but also why “welfare” in particular is the one that can function as covert code, while “assistance for the poor” cannot. Motivated in part by this challenge posed by co-extensional expressions, the linguists Robert Henderson and Elin McCready (2019) develop a refinement of Khoo’s account. Like Khoo, they take covertly coded messages to be pragmatic inferences. But Henderson and McCready focus on a particular kind of inference—one that relies, not on beliefs about the extension of a given expression, but rather on beliefs about what kinds of speakers use that expression.⁵⁶ To better

⁵⁶ Though Khoo notes the possibility that hearers may infer certain coded messages on the basis of beliefs about “the kind of people who use [a given] expression” (2017, p. 48 n. 43), he does not elaborate on that mechanism in this work. He does, however, say more about such beliefs in more recent work on code words (Khoo 2021, p. 153), where he describes such inferences as involving “meta-linguistic bridge principles”—i.e., beliefs about the pattern of usage of a given expression.

capture the latter, Henderson and McCready introduce the notion of *persona*. *Personae* are, roughly, the social identities that hearers associate with a linguistic expression—and not with its semantic content or extension—in virtue of who uses that expression.

On Henderson and McCready’s proposal, the inferences posited by Khoo are triggered only when the speaker utters an expression which the hearer associates with a certain persona. The inference from “inner city” to a racist message is thus mediated by an attribution of a certain persona (e.g., “cryptoracist”) to the speaker (p. 5). Like Khoo’s original account, this refinement encompasses unintentional covertly coded speech. After all, a speaker might be unaware of the persona their hearers associate with a certain expression, and thus unintentionally evince a certain persona. Henderson and McCready also refine the explanation of the in-group/out-group split: only some hearers “hear” the dog whistle because the relevant pragmatic enrichments are available only to hearers who are aware that a certain expression signals a certain persona.

But most importantly, Henderson and McCready offer an explanation for why different co-extensional expressions, such as “welfare” and “assistance to the poor,” don’t necessarily produce the same effects: since *personae* are associated with choice in expression, different expressions with the same content needn’t both signal the same persona. Henderson and McCready point out that such associations between expressions and *personae* are not unusual: “Certain groups of people speak a certain way, and any variation, including lexical choices can signal group membership, a familiar point in sociolinguistic theory” (p. 5). Nor is there anything extraordinary about the fact that competent language users don’t share all the same persona-

associations, since such associations grow out of exposure to how various social groups use language—exposure which of course varies from person to person.⁵⁷

Henderson and McCready do go further than Khoo toward explaining why co-extensional expressions like “welfare” and “assistance for the poor” don’t produce the same coded effects. But for the notion of persona to explain this phenomenon, we’d need to say more about *why* hearers associate certain expressions with certain personae. And we’d also need to explain why millions of hearers, in different places and times, converge on associating the same persona with the same expression. After all, it can’t simply be a coincidental alignment of individual language users’ beliefs. This is clearest for expressions whose coded usage is well established. For example, an utterance of “law and order” today can send a coded message even to hearers who’ve never heard the phrase used by blatant racists (as during the Jim Crow era, for example). How did these hearers come to associate “law and order” with racism, if every utterance of it they’ve heard has come with plausible deniability insulating the speaker from that racist persona?

To answer this question, the notion of a persona must be understood as a structural, relational feature of the linguistic practice itself, and not as a generalization about the experiences of an individual speaker considered in isolation. In other words, familiarity with the coded practice needn’t result from individual speakers’ each reasoning, “Speaker A is a racist, and A used ‘law and order,’ B is a racist and used ‘law and order,’ ... so ‘law and order’ must be used by racists.” Rather, we can also learn that there is such a practice, for example via the reports of others or by careful study of history. Crucially, that history—and in particular the

⁵⁷ Thus, as in my account, the in-group, per Henderson and McCready’s account, includes anyone familiar with how, e.g., racists talk, regardless of whether they would talk that way or endorse that talk or its message.

history of usage—can serve as the very evidence of racism. That is, in learning about the history of “law and order” in American politics, we learn that it’s a phrase used in the service of racist laws, policies, and practices. Equipped with this familiarity with the history of the practice, contemporary English speakers can associate the use of the phrase with racism, even if they lack the independent evidence that would substantiate each of the claims in the sequence, “Speaker A is a racist, and A used ‘law and order,’ B is a racist and used ‘law and order,’ ...”

Earlier, I said that an account of coded speech should explain the difference in effects between “welfare” and “assistance for the poor,” and why “welfare” can function as covert code while “assistance for the poor” cannot. By identifying the mechanisms internal to hearers’ minds, which could produce such effects, Khoo (2017) and Henderson and McCready (2019) offer one part of the explanation. But, as I’ve argued, there are still some details missing. To fill in these details, we’d study the history of our practices of using “welfare” and “assistance to the poor,” and identify the relations between these practices. We’d find, besides the widely shared use of “welfare,” a practice in which “welfare” sends racist messages in certain contexts, a practice with which some but not all competent users of “welfare” are familiar. It’s this familiarity that the in-group brings to their interpretation of an utterance, and that therefore shapes the different messages “welfare” sends to the in- and out-groups, or the different effects an utterance of “welfare” produces on the in- and out-groups. Assuming we don’t find such a pair of coded practices for “assistance to the poor,” we’ll be able to explain why “welfare” and “assistance to the poor” produce different effects, even as they seem to refer to the same thing. Only a practice-focused approach can deliver these missing details. In sum, even if we conceive of covertly coded speech in terms of hearers’ inferences, we’ll need to take the practice-focused approach, if

we're to explain the conformity among many hearers' inferences, across various contexts and generations.

6. Conclusion

In this paper I've identified a certain phenomenon appearing in our actual speech practices, covertly coded speech, in terms of three distinctive features—differential effects, hiddenness, and recurrence. As with any communication in a natural language, covertly coded speech is only possible because there's something a speaker and hearer are both part of—a linguistic practice, an established usage learned from others and shaped by others' past performances. I've thus identified the enabling conditions for such speech in terms of certain pre-existing practices. Covertly coded speech, I've argued, exploits a pair of practices that involve the same expression and that can be performed in at least some of the same contexts, but that send different messages, and that differ in how widespread or narrow they are. Only in linguistic communities structured in this way, and before an audience that includes some who are and some who aren't familiar with the relevant narrow practice, can a speaker send, whether intentionally or unintentionally, the kind of systematically recurring, covertly coded messages that, for example, “inner city” or “law and order” send.

My second, methodological aim has been to show that studying covertly coded speech in terms of the communal practices involved is crucial. Understanding how covertly coded speech works requires attending to the structure of the linguistic community, and to the histories of and interactions among that community's and its sub-communities' linguistic practices. From a purely individualistic, ahistorical, or utterance-level view, we misunderstand or even overlook how and when covertly coded speech is possible. Crucially, the explanatory purchase that comes

with a practice-focused approach does so precisely via attention to the historical, social, and political details—that is, not by taking a synchronic view that assumes a single, monolithic linguistic practice, but rather by focusing on the relations between and structures of multiple interactive practices, practices that are temporally extended, socio-politically shaped, historically and materially embedded.

Using this tool does involve shifting our focus away from particular utterances and individual speakers, but this doesn't make the practice-focused approach incompatible with other approaches to meaning and communication. Focusing on practices is thus a way to extend our existing theories, not replace them. A variety of views of communication can and should make room for the notion of practices and attend to the structures and histories of our linguistic practices. Indeed, even individualist approaches that focus on speakers' communicative intentions or hearers' beliefs and inferences must ultimately appeal to facts about communities' practices, in order to identify the enabling conditions of covertly coded speech and explain the distinctive features. I've thus considered, via the examples discussed in section 5, how the practice-focused approach can be employed within frameworks that ordinarily focus only on individual utterances or individual speakers. What our existing theories stand to gain from the practice-focused approach is an added tool in the philosophy of language toolkit—another mechanism for describing and understanding how we communicate, send messages, get our hearers to believe (or at least entertain) something.

Of course, covertly coded speech isn't just theoretically significant, but also socially and politically significant, for covertly coded speech allows speakers to plausibly deny having sent a certain message. Though this paper has not focused on the sociopolitical significance, it's clear that covertly coded speech is a dangerous tool in the hands of politicians who must build

coalitions to be elected. Their use of covertly coded speech undermines the quality of our public discourse and trust in our public figures, by excluding some members of a political candidate's audience from being given as full and complete an access to that candidate's positions as other members. But *not* using covertly coded speech is liable to preclude a candidate from winning their races against candidates who do use such speech, so it's not enough to simply expect or demand that candidates abstain from covertly coded speech. As I've argued, covertly coded speech relies on audience familiarity with certain practices. Without our competence in these practices, politicians couldn't manipulate us with covertly coded speech. If indeed we find covertly coded speech harmful, then we as audiences have to do our part in recognizing, repudiating, and desisting from the covertly coded practices we may have unknowingly mastered.

Acknowledgements

I would like to give special thanks to David Beaver for his unwavering support and his unique blend of encouragement and challenge, without which this work wouldn't dare be what it is. My sincerest thanks to Jason Stanley, Hans Kamp, Josh Dever, Ray Buchanan, Mark Sainsbury, and David Sosa for many generative discussions and helpful comments. I would also like to thank Brendan Learnihan-Sylvester and Alex Rausch for their feedback on various stages of this work from its earliest forms. Thank you to Laurenz Casser, whose support has been essential during the most challenging stages. I'm grateful to the organizers and audiences of the UT Austin Graduate Colloquium in 2018 and the University of Pittsburgh-based, online Words Workshop in 2020 for their questions and comments, especially Michael Barnes, Laura Caponetto, Nikki Ernst, Megan Hyska, Jonathan Ichikawa, Quill Kukla, Maxime Lepoutre, Sally

McConnell-Ginet, Jennifer Saul, and Bronwyn Stippa. Thank you also to Roy Sorensen and the participants of his Social Epistemology seminar in spring 2021 for stimulating conversation.

Finally, I'd like to thank three anonymous reviewers for their detailed engagement and constructive feedback, which markedly improved the work.

References

- Albertson, B. (2015). Dog-whistle politics: Multivocal communication and religious appeals. *Political Behavior*, 37(1), 3-26.
- Alexander, M. (2012). *The new Jim Crow: Mass incarceration in the age of colorblindness*. The New Press.
- Anderson, L. (2018). Calling, addressing, and appropriation. In D. Sosa (Ed.), *Bad words* (pp. 6-28). Oxford University Press.
- Asher, N. & Lascarides, A. (2001). Indirect speech acts. *Synthese*, 128(1), 183-228.
- Austin, J. L. (1962). *How to do things with words*. Oxford University Press.
- Beaver, D. & Stanley, J. (Forthcoming). *The politics of language: An essay in non-ideal theory*. Princeton University Press.
- Blackmon, D. A. (2009). *Slavery by another name: The re-enslavement of Black Americans from the Civil War to World War II*. Anchor, 2009.
- Bourdieu, P. (1977). *Outline of a theory of practice* (R. Nice, Trans.). Cambridge University Press. (Original work published 1972).
- Brandom, R. (1994). *Making it explicit: Reasoning, representing, and discursive commitment*. Harvard University Press.
- Burge, T. (1979). Individualism and the mental. *Midwest Studies in Philosophy*, 6, 73-121.
- Camp, E. (2018). Insinuation, common ground, and the conversational record. In D. Fogal, D. Harris, & M. Moss (Eds.), *New work on speech acts* (pp. 40-66). Oxford University Press.
- Clark, H. H. (1992). *Arenas of language use*. University of Chicago Press.
- Dewey, J. (1925). *Experience and nature*. Open Court.

- Drakulich, K., Wozniak, K. H., Hagan, J., & Johnson, D. (2020). Race and policing in the 2016 presidential election: Black Lives Matter, the police, and dog whistle politics. *Criminology*, 58, 370-402.
- Eckert, P., & McConnell-Ginet, S. (1992). Think practically and look locally: Language and gender as community-based practice. *Annual Review of Anthropology*, 21(1), 461-488.
- Giddens, A. (1984). *The constitution of society: Outline of the theory of structuration*. Cambridge: Polity Press.
- Gilbert, M. (1989). *On social facts*. Princeton University Press.
- Grice, H. P. (1957). Meaning. *The Philosophical Review*, 66, 377-88.
- Haney López, I. (2014). *Dog whistle politics: How coded racial appeals have reinvented racism and wrecked the middle class*. Oxford University Press.
- Hanna, P., & Harrison, B. (2004). *Word and world: Practice and the foundations of language*. Cambridge University Press.
- Haslanger, S. (2016). What is a (social) structural explanation? *Philosophical Studies*, 173(1), 113-130.
- Haslanger, S. (2018). What is a social practice? *Royal Institute of Philosophy Supplements*, 82, 231-247.
- Henderson, R., & McCready, E. (2019). Dogwhistles and the at-issue/non-at-issue distinction. In D. Gutzmann & K. Turgay, (Eds.), *Secondary content: The semantics and pragmatics of side issues* (pp. 222–45). Brill.
- Hurwitz, J., & Peffley, M. (2005). Playing the Race Card in the Post–Willie Horton Era: The Impact of Racialized Code Words on Support for Punitive Crime Policy. *Public Opinion Quarterly*, 69(1), 99-112.
- Hess, L. (2020). Practices of slur use. *Grazer Philosophische Studien*, 97, 86-105.
- Karttunen, L. (1974). Presuppositions and linguistic context. *Theoretical Linguistics*, 1, 181-194.
- Khoo, J. (2017). Code words in political discourse. *Philosophical Topics*, 45(2), 33-64.
- Khoo, J. (2021). Code words. In J. Khoo & R. Sterken, (Eds.), *The Routledge handbook of social and political philosophy of language* (pp. 147-160). New York: Routledge.
- Kripke, S. (1982). *Wittgenstein on rules and private language*. Harvard University Press.
- Kukla, R. (2014). Performative force, convention, and discursive injustice. *Hypatia*, 29(2), 440-457.
- Kukla, R., & Lance, M. N. (2009). *‘Yo!’ and ‘Lo!’: The pragmatic topography of the space of reasons*. Harvard University Press.

- Lance, M. N., & O'Leary-Hawthorne, J. (1997). *The grammar of meaning: Normativity and semantic discourse*. Cambridge University Press.
- Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge University Press.
- Lepore, E., & Stone, M. (2014). *Imagination and convention: Distinguishing grammar and inference in language*. Oxford University Press.
- Lewis, D. (1969). *Convention: A philosophical study*. Wiley-Blackwell.
- McConnell-Ginet, S. (2011). *Gender, sexuality, and meaning: Linguistic practice and politics*. Oxford University Press.
- McConnell-Ginet, S. (2020). *Words matter: Meaning and power*. Cambridge University Press.
- McMillan, K. (2018). *The constitution of social practices*. Routledge.
- Mead, G. H. (1934). *Mind, self, and society: From the standpoint of a social behaviorist*. University of Chicago Press.
- Mendelberg, T. (2001). *The race card: Campaign strategy, implicit messages, and the norm of equality*. Princeton University Press.
- Putnam, H. (1975). The Meaning of "meaning." *Minnesota Studies in the Philosophy of Science*, 7, 131-193.
- Quaranto, A. & Stanley, J. (2021). Propaganda. In J. Khoo & R. K. Sterken (Eds.), *The Routledge handbook of social and political philosophy of language* (pp. 125-146). Routledge.
- Roberts, C. (2018). Speech acts in discourse context. In D. Fogal, D. Harris, & M. Moss (Eds.), *New work on speech acts* (pp. 317-359). Oxford University Press
- Rouse, J. (2015). *Articulating the world: Conceptual understanding and the scientific image*. University of Chicago Press.
- Saul, J. (2018). Dogwhistles, political manipulation, and philosophy of language. In D. Fogal, D. Harris, & M. Moss (Eds.), *New work on speech acts* (pp. 360-383). Oxford University Press.
- Saul, J, and Drainville, R. (Forthcoming). Visual and Linguistic Dogwhistles. In L. Anderson & E. Lepore (Eds.), *Oxford handbook of applied philosophy of language*. Oxford University Press.
- Sbisà, M. (2002). Speech acts in context. *Language & Communication*, 22(4), 421-436.
- Schatzki, T. (1996). *Social practices: A Wittgensteinian approach to human activity and the social*. Cambridge University Press.
- Schiffer, S. (1972). *Meaning*. Clarendon Press.

- Searle, J. (1995). *The construction of social reality*. The Free Press.
- Smith, T. W. (1987). That which we call welfare by any other name would smell sweeter: An analysis of the impact of question wording on response patterns. *Public Opinion Quarterly*, 51(1), 75-83.
- Stalnaker, R. (1974). Pragmatic presuppositions. In M. Munitz & P. Unger (Eds.), *Semantics and philosophy* (pp. 197–214). New York University Press.
- Stalnaker, R. (1978). Assertion. In P. Cole (Ed.), *Syntax and semantics, vol. 9: Pragmatics* (pp. 315-332). Academic Press.
- Stanley, J. (2015). *How propaganda works*. Princeton University Press.
- Tirrell, L. (1999). Derogatory terms: Racism, sexism, and the inferential role theory of meaning. In C. Hendricks and K. Oliver (Eds.), *Language and liberation: Feminism, philosophy, and language* (pp. 41-79). State University of New York Press.
- Tirrell, L. (2012). Genocidal language games. In I. Maitra & M. K. McGowan (Eds.), *Speech and harm: Controversies over free speech* (pp. 174-221). Oxford University Press.
- Tuomela, R. (2002). *The philosophy of social practices: A collective acceptance view*. Cambridge University Press.
- Winter, N. (2006). Beyond welfare: Framing and the racialization of white opinion on Social Security. *American Journal of Political Science*, 50(2), 400-420.
- Wittgenstein, L. (1953). *Philosophical investigations* (G. E. M. Anscombe, Trans.). Macmillan.
- York, Tony, & MacAlister, D. (2015). *Electronic security system integration*. Elsevier.