

Sociophonetics, semantics, and intention¹

ERIC K. ACTON

Eastern Michigan University

Campbell-Kibler (2008, 2009) observes that the role of speaker-intention seems to differ in the meanings of primary interest in variationist sociolinguistics on one hand and semantics and pragmatics on the other. Taking this observation as its point of departure, the central goal of the present work is to clarify the nature of intention-attribution in general and, at the same time, the nature of these two types of meaning. I submit general principles by which observers determine whether to attribute a particular intention to an agent—principles grounded in observers' estimation of the agent's beliefs, preferences, and assessment of alternative actions. These principles and the attendant discussion clarify the role of alternatives, common ground, and perceptions of naturalness in intention-attribution, illuminate public discourses about agents' intentions, point to challenges for game-theoretic models of interpretation that assume cooperativity, and elucidate the nature of the types of meaning of interest. Examining the role of intention vis-à-vis findings and insights from variationist research and the formally explicit game-theoretic models just mentioned foregrounds important differences and similarities between the two types of meaning of interest and lays bare the contingent nature of all meaning in practice.

Keywords: Intention, social meaning, semantics, pragmatics, indexicality, game theory

1. INTRODUCTION

That buzzing-noise means something. You don't get a buzzing-noise like that, just buzzing and buzzing, without its meaning something.

– Winnie-the-Pooh, upon hearing bees buzzing (Milne 2009 [1926]: 6)

MEANING means different things to different people, not least among those who study meaning for a living. As distinct approaches to the study of language meaning expand and increasingly intersect, there is insight to be gained from closely comparing different types and notions of meaning and their implications. Recent work along these lines has been very clarifying, illuminating the characteristics of and relations between various types of meaning via properties like conventionality, backgroundedness, and projectivity (e.g. Potts 2015, Smith et al. 2010) and performativity and interiority (Eckert 2019).

Grice (1957) identified intention—the focus of the present work—as another concept distinguishing between kinds of meaning. The examples in (1) illustrate.

- (1) (a) Those spots mean measles.
 (b) Those three rings on the bell (of the bus) mean that the 'bus is full.'
 (Grice 1957: 377)

The spots in (1a) 'mean' measles in the sense that they are informative: they tell us that whoever bears the spots has measles. But this doesn't require that anyone means anything BY the spots, or that there's intentionality behind them. In contrast, with (1b) we are invited to imagine a bus driver ringing a bell thrice to signal overtly via convention that the bus is full. Here, someone is presumed to have acted intentionally to communicate something. Indeed, if the ringing seemed unintentional (perhaps due to an involuntary convulsion) the peals would no longer provide reason to believe that the bus was full.

This brings us to the epigraph above. Pooh hears bees buzzing and concludes that the buzzing must mean something. But do the bees mean something BY their buzzing? In a way, it's not such a silly question. After all, one might ask, why should a creature go to the trouble of making sustained, noisy noises if not to

communicate something? But of course buzzing is something bees can't help doing if they want to move themselves about swiftly and independently. Thus, one need not conclude from a bee's buzzing that the bee means anything by it; it may simply be incidental to a goal of locomotion.

Campbell-Kibler (2008, 2009) observes that intention seems to play a different role in the meanings of primary interest in variationist sociolinguistics—meanings based in the stances, traits, and personae associated with and indexed by particular linguistic forms—as compared to the meanings of primary interest in semantics and pragmatics—meanings based in and derived from conventionalized, semantic content. This work takes Campbell-Kibler's observation as its point of departure, with the central goal of clarifying how observers (hearers) determine whether to attribute a particular intention to an agent (speaker) and, relatedly, clarifying the nature of and relationships between the two types of meaning just mentioned.

Section 2 presents what I take to be central principles by which one determines whether to attribute a particular intention to an agent given their action. Sections 3 and 4 examine the role of intention in sociophonetically and semantically based meaning, respectively. I show how the principles presented in Section 2 elucidate the nature of both types of meaning and relate to notions like naturalness, performativity, salience, common ground, and pragmatic inference. As just one example, we will see that the principles help explain why meaning based in the indexical associations of phonetic forms is particularly amenable to being perceived as unintended, especially relative to semantically based meaning. At the same time, the discussion will ultimately make clear that—even with meanings based in an utterance's semantic content—language users must contend with substantial uncertainty concerning speakers' intentions, and they lean on the principles outlined in Section 2 in attempting to discern those intentions.

These points lead into Section 5, where I examine Burnett's (2017, 2019) work on SOCIAL MEANING GAMES (SMGs), which applies game-theoretic pragmatics (e.g. Franke 2009, Frank & Goodman 2012) to the meanings of primary interest in variationist sociolinguistics. Drawing on the discussion from previous

sections, I argue that aspects of such meaning present major challenges for SMGs, particularly given SMGs' assumption of a certain kind of cooperativity between interlocutors. Still, SMGs sharpen our understanding of meaning and interpretation because of the explicitness that formal modeling requires. Indeed, I show that careful consideration of the role of intention vis-à-vis SMGs and previous variationist research helps lay bare the contingent and performance-based nature of all meaning in practice. Section 6 concludes.

2. INTENTION

In regarding an action one may ask whether a particular consequence was intended by the agent. Let us say that for a (potential or actual) consequence of an action to have been *INTENDED* by an agent means that the agent performed the action as they did in part for the purpose of bringing about that consequence. Along these lines, let's say that an action (or aspect thereof) was *INTENTIONAL* if and only if it was committed for the purpose of bringing about one or more of the agent's goals.

These informal definitions will suffice for our purposes. To be clear, I don't mean for an action with an intended consequence to require that the agent can parse out precisely what they did or why they acted as they did in service of that consequence. In attempting to be friendly, for instance, I may do all sorts of things with my posture, voice, etc. that I'm not fully aware of but that I enact purposively towards the goal of appearing friendly. Being in service of an agent's goal, such actions are intentional in the relevant sense, and the goals they serve are likewise intended.

It is worth noting along these same lines that goal-oriented action needn't be consciously orchestrated. As Bargh et al. (2008: 535) put it, '[G]oals can be activated without an act of conscious will [...] and then operate in the absence of conscious guidance to guide cognition and behavior towards the desired end state', noting the separation between structures in the brain related to executive function and conscious awareness. Accordingly, degrees of conscious awareness will not play a central role in the present work. (That said, intuitively people might

assign more responsibility to agents for those aspects of their behavior of which they are consciously aware.)

2.1. *Attributing Intention*

One claim of this work is that meaning based in sociophonetics is generally less likely to be taken to be intended by the speaker than that based in the semantics of morphosyntactic objects. This raises the general question, at the center of this work, of when a potential consequence of an action will be viewed as intended.

To address this question, I begin with the notion of an ACCESSIBLE ALTERNATIVE, which I characterize in (2).

- (2) Given an action α performed by A , an alternative action α' is an ACCESSIBLE ALTERNATIVE to α for A iff A , consciously or subconsciously:
- (a) Knew of the availability of α' prior to performing α ; and
 - (b) Could have performed α'

In essence, an alternative was accessible for agent A if and only if it was an action that A could have performed and was on some level on A 's radar.

Of course, just because an alternative is accessible doesn't mean it's desirable. Actions and their alternatives can come with various potential costs, which may be realized in physical, financial, social, or other terms. Certain actions, for instance, require great effort, making them generally less attractive than other actions. Alternatives also carry potential benefits, which we may think of in terms of the likelihood with which they will bring about desirable outcomes. Such potential costs and benefits determine how attractive a given accessible alternative is to an agent and, as I will discuss shortly, they in turn play an important role in assessing whether a consequence of an action was intended.

I now turn to the principles governing whether an observer takes a consequence of an action to have been intended, given in (3). The principle revolves around how favorably the observer thinks the agent would view the relevant outcome, how likely the observer thinks the agent would think the action was

to effect that outcome, and how the observed action relates to alternatives that appear to have been less or more likely to effect that outcome.

Underlying (3) and the ensuing discussion is the presumption that agents and their observers are rational in the sense of attempting to maximize the net benefits (benefits less costs) of what they do (cf. Horn 2004, Sperber & Wilson 2004), and that this presumption of rationality is common ground among them. Where rationality of this kind isn't assumed, the dynamics outlined below fall apart. It is also worth emphasizing that I do not intend to make any assumptions about the degree of consciousness under which (3) operates.

- (3) **ATTRIBUTING INTENTION.** Suppose *O* observes agent *A* performing action α with potential or actual consequence *c*. Let an *M*-ALTERNATIVE be an alternative action that *O* believes *A* would have thought was accessible and more likely than α to effect *c*; and an *L*-ALTERNATIVE be an alternative that *O* believes *A* would have thought was accessible and less likely than α to effect *c*.

O is more likely to believe *A* intended to effect *c* via α :

- (a) **DESIRABILITY.** The more *O* thinks that *A* would view *c* as desirable
- (b) **EXPECTATION.** The more likely *O* thinks *A* would have thought α was to effect *c*
- (c) ***M*-ALTERNATIVES.** The less *O* thinks *A* believed *A* was forgoing in net benefits by selecting α over *M*-alternatives
- (d) ***L*-ALTERNATIVES.** The more *O* thinks *A* believed *A* was forgoing in net benefits by selecting α over *L*-alternatives

(3a)'s 'Desirability' factor, is perhaps obvious: We are more likely to think *A* intended a particular outcome for *A*'s action the more we think that *A* would desire that outcome. If a footballer misses a shot on goal in an important match, we're generally unlikely to think the miss was intended; typically, players want NOT to miss. On the other hand, if we later learn that the player had bet a large sum on the match against their own team, we have good reason to think the player would

want to miss the goal and thus that they intended to do so. The ‘Expectation’ factor in (3b) is similarly straightforward. The less *A* believes α will bring about *c*, the less sense there is in *A* attempting to effect *c* via α . If, for instance, a two-year-old pulls a fire alarm and we doubt that the child was aware of the likely consequences of doing so, we likewise have reason to doubt the child intended to clear the building of people, etc. But if an adult who presumably has full knowledge of the purpose of fire alarms does so, there is good reason to think that that person indeed intended to clear the building.

(3c) and (3d) are more complex because they involve not just the action itself but also how it relates to alternative actions. (3c), the ‘*M*-alternatives’ factor, says that we are more likely to believe that *A* intended some outcome *c* the less we think *A* thought *A* was forgoing (in net benefits) by selecting α over alternatives that were apparently accessible and more likely than α to effect *c*. As an example, imagine we learn that an individual, Alex, said the word *damn* somewhat frustratedly, and we want to know whether Alex intended to express extreme displeasure, as opposed to, say, only moderate displeasure. Assume that Alex has at other times, especially when very upset, used a more taboo expletive than *damn* to express displeasure. Thinking in terms of (3c), our *c* here is expressing extreme displeasure, and a potential *M*-alternative—that is, an alternative to saying ‘damn’ that we think Alex would regard as accessible and more likely to express extreme displeasure—would be saying that stronger expletive instead.

Now let’s consider two possible settings for Alex’s utterance. In the first, Alex is among friends who frequently swear. In the second, (holding other factors fixed—e.g. Alex’s delivery), Alex is at school teaching a class of kindergartners, where swearing is very much frowned upon. Intuitively speaking, we have greater reason to think Alex is very upset from Alex’s saying *damn* in front of kindergartners than among friends. In the kindergarten scenario, Alex, even if deeply upset, has good reason to avoid a stronger expletive (e.g. the potential of censure), whereas in the other scenario there’s no great social cost to using a

stronger expletive. Putting things again in terms of (3c), we could say that uttering the stronger expletive offers less in net benefits in the kindergarten scenario than in the friends scenario, because in the kindergarten scenario uttering a stronger swear comes at a higher cost. Thus, in accordance with (3c), we are more likely to think Alex intended express extreme displeasure by saying *damn* in the kindergarten scenario, where the net benefits of using a stronger expletive are less.

A related logic underlies the ‘*L*-alternatives’ factor spelled out in (3d), which says that we’re more likely to think *A* intended to effect *c* the more we think *A* would think alternatives less likely to effect *c* had to offer relative to α . To illustrate, consider the hypothetical case of Pat, who is certain to buy one of two houses that are identical except that one is on the east side of the street and the other is directly across from it on the west. Suppose Pat buys the house on the east side, and we are curious as to whether Pat wanted a house on that side in particular. Here, then, our *c* is having Pat’s home be on the east side of the street and the single *L*-alternative—in this case, an alternative action we believe Pat would think was less likely to effect *c*—is buying the house on the west side.

Again, comparing two versions of this scenario will illustrate the substance of (3d). In the first, suppose that the house on the west side cost \$25,000 more than the other. In that case, it’s not clear that Pat wanted to live on the east side in particular—perhaps the lower price was the reason for selecting that house. Now suppose the house on the west side cost \$25,000 LESS than the other. In that case, we have greater reason to think Pat wanted to live on the east side in particular, because Pat could have paid considerably less for an identical house on the other side. Put in terms of (3d), the net benefits of opting for the *L*-alternative—namely, buying the house on the west side of the street—are higher in the second case, where doing so would cost \$25,000 less than Pat’s actual action, than in the first, where it would cost \$25,000 more. Accordingly, as (3d) suggests, we are more likely to think Pat intended to have a home on the east side of the street in particular in the second scenario, where the house on the east side costs more.

As we transition to thinking more specifically about intention in communication, it is worth noting that this *L*-alternatives factor provides grounds for doubting that Pooh's bees intended to communicate something by their buzzing, or even specifically intended to issue a buzzing sound. Assuming the bees wished to move themselves swiftly, any alternative action less likely to involve buzzing (e.g. walking, waiting to be carried by the wind, etc.)—would have little to offer. Hence, according to (3d), the buzzing was quite plausibly not a goal in itself but rather incidental to a goal of locomotion. This is akin to Pat buying a house on the east side of the street when it's \$25,000 cheaper than the alternative—for all we know, it was saving money, not living on the east side in particular, that led to Pat's purchase of that house. The dynamics involving *L*-alternatives in intention-attribution will play a prominent role in the material to come.

To summarize, (3) relates the likelihood of ascribing to an agent *A* the intention to effect some outcome *c* via action α to four factors: (i) *c*'s apparent desirability; (ii) how much we think *A* would expect α to effect *c*; and α 's apparent attractiveness relative to alternative actions that seem (iii) more or (iv) less likely to effect *c*. Thus far, I have discussed actions and outcomes in very broad terms, with examples from multiple domains, to illustrate this principle's generality. Let us now focus more narrowly on a particular kind of action: utterances.

2.2. *Intention and language meaning, broadly construed*

Like actions in general, utterances have potentially many different parts, costs, benefits, and consequences. Among an utterance's possible consequences are effects on a hearer. An utterance may for instance cause a hearer to feel a certain way, recall certain concepts, or change their beliefs. In my view, meaning studies should ultimately encompass a wide range of impacts that utterances have on hearers' emotional and mental states, but I will focus on belief states here.

In talking about meaning and the way the term MEANING is used, it will be useful to have a broad working notion of meaning that's general enough to accommodate the diversity of phenomena labeled meaning in the literature. The

notion described here isn't meant to be absolute; in other contexts, other notions may be warranted. For our purposes, we might think of a form's meaning in the abstract as its potential contribution to suggesting something about the world to a given hearer in a given context. This accords with the idea that a word's semantic meaning is a function that may be composed with other word meanings to form a proposition (e.g. Montague 1970), which, when issued via an utterance, may suggest something about the world. It is likewise accords with the idea that the meaning of an abstract variant of a sociolinguistic variable is an INDEXICAL FIELD (Eckert 2008)—a collection of indexically associated stances, traits, etc.—which, when used, may suggest that the speaker embodies or wishes to mark as relevant some subset of that field. Moving away from forms in the abstract, we might say that, for a given hearer, the meaning, in this very broad sense, of a particular instantiation of a form in practice is whatever contribution it actually makes to suggesting something to that hearer.

This very general notion of meaning-in-practice is stated relative to the hearer and says nothing about intention. Let's now incorporate the perspective of the speaker. Suppose a speaker *S*'s utterance *u* suggests something *p* to a hearer *H*. (4) enumerates three relations that may obtain between *S*, *H*, *u*, and *p* of particular interest herein, related to Grice's (1957) observations.

- (4) UTTERANCES AND SPEAKER INTENTIONS. Let *S* be a speaker whose utterance *u* suggests something *p* to hearer *H*. Among the possible relations between *S*, *u*, *p*, and *H* are:
- (a) *S* had no intention concerning *u* suggesting *p* to *H*.
 - (b) *S* intended *u* to suggest *p* to *H* and intended for *H* NOT to recognize that intention.
 - (c) *S* intended for *u* to suggest *p* to *H* and for *H* to recognize that intention.

As alluded to in Section 1, Grice (1957) distinguishes between two types of meaning: 'natural' and 'nonnatural'. Natural meaning roughly coincides with

(4a), where something provides information without any particular intention behind it. (Grice's discussion suggests something a bit stricter, but this suffices for our purposes.) This is the 'Those spots mean measles' case. In contrast, to say that *S* meant something by some utterance *u* in the 'nonnatural' sense, according to Grice (1957: 385), 'is (roughly) equivalent to "[*S*] intended [*u*] to produce some effect in an audience by means of the recognition of this intention"'.² In keeping with this notion, the vast majority of research in semantics and pragmatics has focused on instances where the relation in (4c) holds (though see e.g. Ariel 2004, Franke et al. 2012). As Campbell-Kibler (2008, 2009), points out, however, very often in sociolinguistic research what a speaker's utterance suggests to the hearer is not assumed to have been intended. Indeed, the remainder of this work makes clear that all of the relations in (4) are possible in practice, whether the source of a meaning is an utterance's semantic content or its phonetic rendering.

2.3. *Looking ahead*

In Sections 3 and 4, I will employ the principle of intention-attribution in (3) and the relations enumerated in (4) to illuminate the nature of both meaning based in the indexical associations of variants of sociolinguistic variables, often referred to as SOCIAL MEANING, and meaning based in the semantic content of morphosyntactic expressions. Three major themes will emerge.

First, utterances can suggest information in diverse ways—some based in indexical associations, some based in semantics; some unintended, some overtly or covertly intended. Second, in accordance with previous sociolinguistic and pragmatic research, language users must contend with vast uncertainty in determining what exactly a speaker intends. Third, (3) furnishes principled insights concerning how language users navigate through that uncertainty and concerning the nature of the two types of meaning of interest.

Sections 3 and 4 focus primarily on indexically and semantically based meaning, respectively, but, for the sake of comparison, both sections will refer to both kinds of meaning at times. Section 3 is organized around the three

intentionality relations in (4). Section 4 features two examples of debates about speakers' intentions—debates which I show are rooted in the principle of intention-attribution in (3). In addition to elucidating the nature of meaning in practice in their own right, these two sections also set the stage for Section 5, where I show that game-theoretic models that only address meaning of the type in (4c) miss much of the crucial complexity and scope of meaning in practice.

3. INTENTIONALITY AND INDEXICAL SOCIAL MEANING

A look the sociolinguistics literature reveals that the notion of 'social meaning' is meant to include meaning that isn't assumed to have been intended by the relevant speaker. Campbell-Kibler (2009: 136), for instance, defines social meaning as 'social content tied in the minds of a given speaker/hearer to a particular piece of linguistic behavior'. A linguistic behavior may therefore be socially meaningful in this sense simply by suggesting something about the social world, irrespective of speakers' intentions. This is consonant with Eckert & McConnell-Ginet's (2013: 490) claim that verbal performances 'come off as something regardless of intention, MEAN SOMETHING [...] because they [...] reiterat[e] what has worked in the past' (emphasis added). (See also Lavandera 1978: 173.)

This is not to say that intention plays no role in social meaning—indeed, speakers as 'stylistic agents' (Eckert 2012: 97–98) looking to send social signals are at the center of meaning-based approaches to sociolinguistics. Rather, the point is that speaker intention is not a PREREQUISITE for social meaning. I will now discuss a diversity of ways in which an utterance or variant may be socially meaningful in the broad sense, taking each of the relations in (4) in turn and highlighting key theoretical points and implications along the way.

3.1. *No intention*

This class of cases aligns with (4a), where something about a speaker *S*'s utterance suggests something *c*, but *S* had no specific intention to suggest *c*. Examples abound in the sociolinguistics literature. Indeed, sociophonetic

perception studies often find, as Campbell-Kibler (2008: 648) puts it, that ‘listeners feel entitled to read qualities into a speaker’s linguistic cues that speakers are unlikely to have included deliberately’. Levon (2014) and Tamminga (2017), for instance, report that speakers’ using particular phonetic variants can suggest to hearers that the speakers are less ‘competent’ and ‘stupid[er]’. Coming off as stupid or incompetent is generally not something one intends to do (with some exceptions, no doubt), just as missing a shot on goal is generally something a footballer doesn’t intend to do. Consistent with (3a)’s desirability factor, then, these meanings are apt to being treated not as something the speaker intended to evince, but rather as revealing something about the speaker’s true nature.

Of course, unintended social meanings need not be undesirable. The main point here is that there are things our speech might signal that we have no specific intention of signaling, whether desirable or not. As just one example, if a speaker *S*, being from the U.S., says something in *S*’s native vernacular, *S*’s accent may suggest that *S* is from the U.S., even if suggesting as much was in no way a motivating factor for *S*’s utterance or its phonetic shape.

3.1.1. *Intention, L-alternatives, and meaning in language*

This brings us to the first key theoretical point of this section, furnished by the *L*-alternatives factor of (3d). (3d) says an observer *O* is more likely to believe that an agent *A* intended to effect some consequence *c* by their action α the more *O* thinks *A* believed that *A* was forgoing by selecting α over *L*-alternatives—alternatives *O* believes *A* would think were accessible and less likely to effect *c*. Regarding the case of language, an important corollary of this principle is given in (5).

(5) ATTRACTIVE *L*-ALTERNATIVES (A corollary of (3d))

Believing that a speaker *S* intended to suggest *p* via an utterance *u* (or aspect of *u*) is:

- (a) Facilitated by believing that *S* forwent an attractive *L*-alternative—i.e., an alternative that *S* would think was accessible, less likely to suggest *p*, and attractive in its net benefits.

(b) Inhibited by believing that *S* had no such alternative available.

To illustrate, let us return to the case of a speaker *S* whose phonology suggests U.S. nationality. Suppose that for all hearer *H* knows, *S* is using *S*'s native phonology and for *S* to speak in a way substantially less likely to suggest U.S. nationality would require greater effort and sound forced. In turn, from *H*'s perspective any such alternative ways of speaking would be rather unattractive to *S* (assuming *S* doesn't want to sound pretentious), requiring more effort with no clear offsetting benefit. In this case, then, there's rather little reason for *H* to think that *S* specifically intended to signal being from the U.S. phonetically. The situation is akin to the case of Pooh's bees—as far as *H* knows, if *S* is to talk without expending unnecessary effort and sounding strange, *S* will evince U.S. nationality simply by speaking. In other words, *S* is 'just buzzing'.

Indeed, though not in exactly these terms, the matter of whether a speaker is 'just buzzing'—just speaking 'naturally'—is often found in discussions of social meaning and intentionality.³ One thing that makes an action attractive for an agent is its being easy to commit, and speaking in a 'natural' manner (however real that notion actually is), intuitively means speaking less effortfully. In keeping with the discussion above, then, if a speaker *S* appears to be speaking naturally and hence relatively effortlessly, the evidence that *S*'s phonetics were intended to serve some special purpose other than just articulating *S*'s utterance diminishes. Contrariwise, if *S* appears to have forgone a more natural, easier way of talking, there's some reason to believe that *S* expended the extra effort towards some goal that *S* believes the more natural alternative wouldn't have served as well.⁴

This very dynamic shows up in Campbell-Kibler's (2008: 648) analysis of how the speaker 'Elizabeth' is evaluated by participants in a study of perceptions of phonetic variants of the English *-ing* suffix: participants interpret Elizabeth as making a 'sociolinguistic move' precisely when she uses a variant that participants judge to be less 'natural' for her. The same basic reasoning about alternatives also helps explain Podesva's (2011) claim that SALIENT variants of sociolinguistic variables are particularly useful for conveying social meaning, where tokens are

salient by being ‘infrequen[t]’ or ‘by exhibiting extreme acoustic values’ (237). Insofar as there is a direct correlation between the frequency and the ease with which one produces a particular form, infrequent forms and forms exhibiting ‘extreme acoustic values’ are presumably more effortful than their counterparts. When a speaker uses such a form, then, there’s reason to believe that the speaker has forgone a less effortful option. If rational, the speaker wouldn’t exert extra effort for no reason, suggesting that the speaker aimed to achieve something that apparently otherwise attractive and less effortful options would be less likely to achieve. One plausible explanation, depending on the context, is that the speaker exerted extra effort in order to convey some social meaning.

Before moving on, it is worth noting that the principle in (5) operates not just vis-à-vis the socioindexicality of phonetic forms but also relative to semantic content. (Indeed, being based in (3), it generalizes from utterances to actions more broadly.) The familiar phenomenon of scalar implicature (e.g. Horn 2004) illustrates. Consider (6) and (7), both evaluated in an otherwise identical context where it is common ground that the respondent is sincere, fully knowledgeable about Terry’s performance, aware of terms like *great*, and prefers to be informative so long as it doesn’t require lying.

(6) Fill in the blank: ‘Terry’s performance was satisfactory.’

(7) Mark the option that applies:

‘Terry’s performance was: unsatisfactory satisfactory.’

The claim here is that though the respondent has given the same answer in both cases, (6) is more likely than (7) to indicate an intention to suggest that Terry’s performance wasn’t great. This Gricean dynamic can be put in terms of the discussion in this section. With (6), the respondent had available alternatives less likely than *satisfactory* to suggest ‘not great’, like *great*. And, given what we know about the respondent, *great* is at least on some dimensions an attractive alternative, being more informative than *satisfactory* without requiring extra articulatory effort. The respondent’s choosing a response more likely to suggest ‘not great’

provides evidence for the possibility that the respondent wished to suggest as much and avoided otherwise apparently attractive alternatives like *great* to serve that goal.⁵ With (7), however, there is no such alternative available, and the corresponding evidence vanishes.

To summarize this section, a corollary of (3d)'s *L*-alternatives factor in intention-attribution, spelled out in (5), is that a speaker is more (less) likely to be taken to have intended to signal *p* by their utterance when it seems there was (wasn't) an attractive alternative less likely to suggest *p* that the speaker forwent. Concerning the phonetic rendering of one's utterance, this means that if a speaker appears to be speaking in a manner in some way unusual for that speaker, the speaker's phonetics are relatively likely to be taken as intended to serve some special purpose other than just articulating their utterance. If the speaker appears to be speaking 'naturally' and hence relatively effortlessly, however, the speaker's phonetics are relatively unlikely to be taken as having some such special purpose.

3.1.2. *Indexical meaning's relative amenability to non-ascription of intention*

The preceding discussion, especially regarding speaking 'naturally', foregrounds another point worthy of emphasis: meaning based in phonetic forms' indexical (or iconic) character is relatively amenable to being perceived as being unintended.⁶ This is because just as bees flying requires rapid vibration of their wings, speaking requires phonetics, which opens the question of whether the phonetic nature of the utterance was just buzzing—that is, whether it was simply in service of attempting to articulate the relevant morphosyntactic objects efficiently—or if it was indeed designed to convey something beyond the semantics of those morphosyntactic objects (and any pragmatic inferences derived therefrom).

The picture is rather different when it comes to the semantics of morphosyntactic objects. Again, as (5) indicates, believing that an aspect of an utterance wasn't intended to suggest something is facilitated by believing that there was no alternative available to the speaker given their other goals and constraints. But words and phrases typically do have accessible alternatives, often including

saying nothing at all. The absurd exchange in (8) illustrates.

- (8) A: The mug is over there under the green coffee table.
 B: All I see there is a book./It looks like it's on top of the coffee table./I only see a brown coffee table.
 A: Yes, that's right.
 B: Then why did you say 'mug'/'under'/it was green?
 A: Oh, I didn't mean anything by it.
 B: ???

Bees' buzzing is incidental to their flying, and my variants of /s/ and /i/ in a given context may simply be in service of articulating the word *see* intelligibly and relatively effortlessly, but it would take very special circumstances for one to observe A's first utterance in (8) and conclude that A used the words *mug*, *under*, and *green* for some reason having nothing to do with evoking their semantic content. Why should A go to the trouble of saying those particular words if A meant nothing by them, especially as part of an ostensibly communicative act? It is indeed generally very strange (though see Section 4.1) to issue a morphosyntactic object, with all of its encoded semantic meaning, with no intention of suggesting something involving that meaning. Morphosyntactic objects, then, are generally harder to write off as just buzzing. Rather, they're generally taken to be intended to suggest something in their own right.

The notion that indexical meaning is relatively amenable to being read as not intended figures importantly in cases of indexical meaning bearing the relation in (4b), to which I now turn.

3.2. *Intention with intention for non-recognition of intention*

This class of cases aligns with (4b), where the speaker intends to suggest something without that intention being recognized. Such cases in a sense depend on the class of cases above: the speaker hopes that the hearer will misidentify the relevant meaning as unintended. I may, for instance, want to sound smart,

cool, etc., but emphatically not want you to recognize as much, as recognition may undermine my intention. Campbell-Kibler (2009: 137) puts it this way: ‘[Some meanings,] such as jaded, lose their meaning when recognized as intentional’. This dynamic illustrated clearly in Campbell-Kibler’s (2008) analysis of participants’ evaluation of speaker ‘Valerie’, whose ‘use of *-ing* [as opposed to *-in*’ ...] successfully means intelligent to exactly those listeners who do not perceive that move as intentional. Those who think she intends “intelligence” by her social cues react by seeing her as less intelligent’ (655).

Thus, successfully suggesting certain things sometimes REQUIRES that the relevant intention go unrecognized. And this stands to reason. If we are trying to discern the kind of person someone is or the state they’re in, we want to get it right, and the appearance of authenticity can be crucial. Being perceived as authentic often depends on being perceived as not going out of one’s way to convince others of something about oneself—that is, it often depends on the observer’s believing that one is just letting one’s true self or state show. A child’s cough, for instance, will only earn them a day home from school if it is believably unintentional.

As just established in Section 3.1.2, meaning based in the indexicality of phonetic forms is particularly amenable to being perceived as unintended. In turn, it is likewise relatively useful for meanings that speakers intend but want to be perceived as unintended.

3.2.1. The performativity of sociophonetics and consequences for reliability

This talk of authenticity brings up the notion of PERFORMATIVITY, which has multiple related senses (e.g. Austin 1962, Butler 1993). Here, in the spirit of Eckert (2019), I’m focused on the degree to which the information suggested by an utterance—and one’s evaluation of the truth of that information—depends on the delivery, that is, the performance, of the utterance itself (along with other aspects of the speaker’s behavior). As Eckert observes, sociophonetic meaning is primarily performative. One’s inexplicit claims to coolness or intelligence, for instance, depend entirely on one’s success in consistently enacting those

traits, which includes any phonetic work towards those ends. Likewise, if one's phonetics make one sound angry in a particular context, then, for all an observer knows, one is indeed angry; there is no objective, outside metric against which to test our evaluations. Given the performative and (inter)subjective nature of sociophonetic meaning, it is thus very often difficult to decisively verify much of what's suggested by the phonetic character of an utterance.

In contrast, the semantic content of one's utterance and the evaluation of the truth of much of that content generally doesn't depend so heavily on delivery. The semantics and truth of the sentences *Springfield is the capital of Illinois* and *Six squared is 36*, for instance, can be determined independently of how they are said.

In this way, in many cases there is a better chance of establishing whether the semantic content of an utterance is true than whether the social information suggested by its phonetics reflects the speaker's 'true' self or stance. Moreover, as suggested by Section 3.1.2, one may in many cases plausibly deny having intended to signal some social meaning via one's phonetics, but it is generally difficult to deny committing oneself to the semantic content of one's utterance (outside of contexts like sarcasm and involuntary or reported speech).

Taking all of this together, then, the semantic content of an utterance generally provides more decisive information about a speaker's tendency to willingly suggest (un)reliable information to their hearers than does its phonetic character. With sociophonetically based meaning, a speaker may persist in acting as though an inferred meaning that has been called into question is indeed 'true', or simply deny ever having intended to suggest it at all.

I will return to the issue of performativity and reliability in the discussion of game-theoretic approaches to meaning in Section 5. At any rate, it should by now be clear that cases where an utterance suggests something but the speaker either didn't intend it to or intended for their intention to go unrecognized are by no means at the periphery of sociophonetic meaning. There are, however, also cases where speakers intend to suggest something via their phonetics and intend for that suggestion to be recognized. Such cases are the focus of the next section.

3.3. *Intention with intention for intention-recognition*

This class of cases lines up with (4c). Podesva's (2011) research on high-rising terminals (HRT) provides an example. Analyzing previous research and his own data on HRT, Podesva (2011: 245) suggests that across various uses, HRT 'serves as a politeness strategy enabling the speaker to express concern for the hearer'. Expressing concern for one's hearer is generally compatible with intending to have one's intention be recognized. King's (2018) research on language in a sex-education class, where a student selectively employs 'Hip Hop styling' prosody to overtly index dominance, provides another example of using phonetics to convey social meaning and intending to have that intention recognized.

One can find similar dynamics at the segmental level. Acton (forthcoming) discusses a case involving a speaker *S* who is a highly-educated, friendly, scarce user of the *-in'* form of the *-ing* suffix, where this information about *S* is common ground between *S* and the hearer *H*. Asked about *S*'s weekend plans, *S* says, 'I'm goin' fishin'!'. Given that *S*'s phonetic tendencies are common ground between *S* and *H* here, it's easy to imagine that *S* used *-in'* instead of *-ing* hoping that *H* would recognize that *S* used what for *S* is a rare and thus presumably relatively costly form in order to suggest something. 'If *S* didn't intend to suggest anything by their phonetics', *S* might count on *H* to ask, 'why didn't *S* use the "cheaper" form that we both know *S* uses more frequently?' In that case, *S* both attempts to signal something and intends for that attempt to be recognized.

H's conclusions about what exactly *S* was trying to signal would depend on multiple factors. As Acton (forthcoming) notes, one possible conclusion would be that *S* invoked the trait-based meanings of *-in'* (e.g. suggesting the speaker is easygoing) not to say something general about *S*, but to say something more local—for instance, that the fishing event will be a casual affair.⁷ In any case, what's important here is that *S* may well intend to suggest something to *H* AND have *H* recognize that intention.

Before turning to the role of intention in meaning at the morphosyntactic level, I wish to highlight the general importance of common ground in cases where the

speaker intends for their communicative intention to be recognized. As has been referenced several times, (5) suggests that a hearer *H*'s believing that a speaker *S* intended to suggest something *p* is facilitated by *H*'s believing that *S* forwent an attractive *L*-alternative. The related claim here is that *H*'s believing that *S* intended for *S*'s intention to suggest *p* to be recognized is facilitated by it being common ground (in the sense of e.g. Stalnaker 2002) between *S* and *H* that *S* indeed had access to an attractive *L*-alternative.

The *goin' fishin'* example provides an illustration of this very general dynamic. Here, *S* relies on *H*'s knowledge that, for *S*, *-ing* is an attractive *L*-alternative to *-in'*, being *S*'s more frequent variant. Because of *H*'s knowledge, *S* can count on *H* to understand that by using *-in'*, *S* is 'going out of the way' phonetically—potentially for some special purpose. In contrast, if the scenario were such that *S* believed *H* had no knowledge of *S*'s phonetic tendencies, *S* would in turn have less reason to expect *H* to conclude that *S* was going out of the way phonetically to communicate something special by using *-in'* rather than *-ing*.

In this way, one's knowledge (or lack thereof) of speakers' repertoire and tendencies—that is, what options are available to them and the relative costs of production of each—has a crucial impact on one's interpretation of their utterances. It makes good sense, then, that studies such as Podesva (2011) and Podesva et al. (2015) offer speaker-specific analyses of the social meanings of phonetic forms that incorporate facts about individual speakers' usage patterns. Very generally, such facts play a pivotal role in discerning what speakers intend and, in turn, what we take away from their utterances.

3.4. *Summarizing*

In this section, we have seen that meaning based on the indexicality of variants of linguistic variables can occur with any of the three intention relations of interest in (4). A phonetic aspect of a speaker's utterance may suggest something despite the speaker having no intention for it to do so, or the speaker may intend for it to suggest something hoping that that intention won't be recognized or that it

will. Nor is it surprising that the first two classes of intention relations receive so much attention in sociolinguistic research. For one, though it's often unclear that a speaker meant anything by the phonetic character of their utterance, that phonetic character is virtually always potentially meaningful in the broad sense of suggesting something to the hearer. Secondly, signaling certain things depends on not appearing to be trying to do so (Section 3.2.1).

I also showed in this section that, though speakers' sociophonetic intentions cannot be known with certainty, (3) provides principled insights into how language users attempt to discern such intentions (or lack thereof). Playing a starring role in that regard was (5)—a corollary of (3d)'s *L*-alternatives factor—which states that believing a speaker intended to signal something *p* via their utterance is facilitated by believing the speaker forwent an attractive *L*-alternative. (I further showed that a hearer *H*'s believing that the speaker *S* intended for their intention to be recognized is facilitated by it being common ground between *S* and *H* that *S* forwent such an alternative (Section 3.3).) (5) squares with the findings of previous sociophonetic perception studies, helps derive Podesva's (2011) claim that 'salient' variants of linguistic variables are particularly useful for conveying social meaning, and explains why, relative to meaning based in the semantics of morphosyntactic objects, meaning based in the indexical character of phonetic forms is particularly amenable to being perceived as unintended.

On this last point, though meaning based in the semantics of morphosyntactic expressions may generally be less amenable to being perceived as having no communicative intention behind it at all, it is not entirely unamenable to such perceptions. Moreover, what exactly a speaker intends to convey via their morphosyntax is very often open to question. Thus, as I will show in the next section, meaning of this stripe—like indexically based meaning—is profoundly underdetermined, and determining what speakers intend by the semantics of their utterances is governed by the same principles of intention-attribution in (3) applied in the foregoing discussion.

4. INTENTION AND SEMANTICALLY BASED MEANING

In this section I will focus on the role of intention in semantically based meaning and how our inferences about speakers' intentions relative to the semantics of their utterances are guided by the principles developed in Sections 2–3. I begin with cases where one might indeed doubt that a speaker intended to suggest anything at all by the morphosyntax of their utterance. I then turn to a pair of illustrative examples where a speaker's intention is contested, showing how the debate about the speaker's intentions is rooted in the principles developed above.

4.1. *Morphosyntactic objects without speaker intention*

There are certainly some cases where individuals issue a word or more without themselves intending to suggest anything by it. Involuntarily speech (e.g. sleep-talking) presumably falls under this category. Verbal 'tics' may fit under this rubric in some cases, too, though it should be noted that research on filled pauses, for instance, has revealed correlations between particular filled pauses (e.g. *um* vs. *uh*) and particular discourse situations (e.g. Clark & Fox Tree 2002), suggesting that the forms are, at minimum, differentially suited for particular intentional purposes. Then there are cases where one speaks on another's behalf (rather than speaking to communicate something one independently wishes to convey), which Levinson (1988) argues have received insufficient attention. Du Bois (1993) and Gaenszle (2016), for instance, discuss the case of divination, where the speaker may simply be a channel for a deity's message.

Expressives, too, present an interesting case. On some analyses, felicitously using expressives doesn't strictly require an intention to suggest anything (see e.g. Bach 2006), though it may REVEAL something about a speaker's emotional state. It certainly seems that we sometimes utter expressives not to suggest anything to anyone but for catharsis (Wharton 2016)—one may exclaim, 'Damn!' upon learning bad news with no one else around, not to communicate but to release anger. But as Blakemore (2013) and Wharton (2016) point out, expressives can also be used with an intention to suggest something, even ostensibly—for

example, to make clear to others one's emotional reaction to something.

Aside from involuntary/compulsory or cathartic speech and channeling the speech of others, however, it is rather rare for one to speak and be perceived as having no intention to suggest anything by the semantics of one's utterance. The remainder of this section concerns cases where a speaker does indeed appear to intend something by their utterance, but what exactly was intended is the subject of debate. While cases where speakers intend to have their intentions recognized are the bread and butter of semantic and pragmatic research (though see e.g. Franke et al. 2012), in practice there is often uncertainty even at the morphosyntactic level about what a speaker intends to suggest. The following examples illustrate how reasoning about what in particular a speaker intended to suggest is grounded in the theory of intention-attribution outlined in (3).

4.2. Example: Ilhan Omar's statements on Israel and its supporters

In this example it is clear that the speaker meant something by her utterances (and something negative at that), but there was controversy concerning whether she specifically intended for her words to have anti-Semitic force. The case involves U.S. congresswoman Ilhan Omar, who in 2018 and 2019 drew considerable attention for statements she made criticizing Israel and its supporters. The most controversial statements are provided here in (9). Space limitations preclude a comprehensive treatment of these statements and their contexts, nor do I intend to endorse a particular interpretation of them. My focus here is on how discussions of Omar's intentions center on the principle in (3).

- (9) (a) **'Israel has hypnotized the world**, may Allah awaken the people and help them see the evil doings of Israel.' (Twitter, 2012 (deleted))
- (b) '[The level of U.S. Congressional support of Israel is] **all about the Benjamins baby!**' (Twitter, 10 Feb 2019) [*Benjamins* is a slang term for \$100 bills]
- (c) 'I want to talk about the political influence in this country that says it is okay for people to **push for allegiance to a foreign country.**'⁸

(Public forum, Washington, D.C., 27 Feb 2019)

Many commentators pointed out that Omar's statements in (9) bear connections to anti-Semitic stereotypes (e.g. Weiss 2019, Beauchamp 2019), and there was no shortage of speculation concerning whether Omar intended for her statements to bear anti-Semitic force. 'Omar, I suspect, knows exactly what she is doing,' wrote *New York Times* columnist Bret Stephens in a 7 March 2019 opinion piece on Omar's statements, and many readers commented for or against this claim (Stephens 2019: n.p). A few illustrative examples are provided in (10).

- (10) (a) User DO5: 'Representative Omar [...] made anti-Semitic comments. No one naturally states, without some previous introduction, the specific, anti-Semitic tropes developed over centuries.'
- (b) User Sarah: '[W]e can't just say it's lack of "tact." After several instances of stepping in it, her approach is either intentional or really, really lazy.'
- (c) User Sedanchair: '@DO5 It seems there is a never-ending list of increasingly obscure tropes to draw from [...] for use against any critic of Israel.'
- (d) User Metastasis: '[...] criticism of a government is not criticism of its people [...]

(10a) suggests that it's not credible that Omar invoked three different anti-Semitic tropes simply by chance and hence must have known her words would have anti-Semitic force. This argument links to the Expectation factor of (3b), which says that we're more likely to think someone intended an outcome *c* the more we think they expected their action to effect that *c*. In terms of our metaphor from Section 2.1, (10a) and Stephens suggest that Omar's actions are not akin to an unwitting two-year-old pulling a fire alarm. (10b) gives a similar assessment, suggesting that, at best, Omar has been indefensibly careless in wording her criticisms of Israel. (10c), on the other hand, offers a counterargument to (10a–b), appealing to (3d)'s *L*-alternatives factor (though not in so many words): according

to (10c), there are no alternative utterances critical of Israel and its supporters but less likely than Omar's statements to be linked to anti-Semitic tropes. In other words, according to (10c), there is simply no way for Omar to criticize Israel without being accused of evoking anti-Semitic tropes, even if the latter is not among her goals.

(10d) defends Omar against claims of anti-Semitism by suggesting that if she had wanted to criticize Jewish or Israeli people in general she could have done so explicitly, rather than talking about Israel as a state. Linked to (3c)'s *M*-alternatives factor, the reasoning here is that among the alternatives accessible to Omar were utterances more likely to have clear anti-Semitic force, but Omar opted against such alternatives, suggesting that expressing anti-Semitism was not her intention. In terms of our earlier metaphor for this factor, (10d) is in effect saying that we have a situation like a speaker saying *damn* among friends: if the speaker really wanted to express extreme displeasure, they could have used a more taboo expletive. On the other hand, (3c)'s *M*-alternatives factor also provides a basis for a counterargument because it takes into account the costs of alternatives. So, while there were accessible alternatives more likely to express unequivocal anti-Semitism, such alternatives may have been so socially costly—perhaps costing Omar her seat in Congress—that, on balance, they offered Omar relatively little even if she did wish to signal anti-Semitism. Looking beyond this particular case, (3c) is what permits one to think that although a speaker forwent more direct or potent routes to effecting some outcome *c*, they may still have intended to effect *c*, avoiding the alternatives simply because they were too costly. This recalls the case of the speaker saying *damn* while teaching a class of kindergartners—the speaker may have forgone a stronger expletive not for lack of intense displeasure but because stronger language might have led to censure.

Stephens' headline and the comments in (10) thus all appeal to the factors in (3), making direct or indirect connections to what Omar might reasonably have expected her utterances to convey or to the properties of alternative utterances more or less likely to convey anti-Semitism.

4.3. *On virtue signaling*

There is no shortage of other examples where a speaker's intentions in issuing some semantic content are underdetermined by that content and potentially covert. In discussing social meaning, I mentioned that speakers may use specific phonetics intending to suggest something about themselves, hoping for that intention to go unrecognized—for example, to sound cool or intelligent. The same goes for semantic meaning. Bach (2012: 52), for instance, notes that one might say self-deprecating things in order to appear modest, hoping for that intention to go unrecognized because recognition of the intention 'may vitiate it'. Such is the nature of so-called 'virtue signaling', where one expresses or does something supposedly virtuous primarily to appear virtuous. Accusations of virtue signaling abound on the Web, as in the tweet in (11) from 4 August 2019. Here the user, whose other tweets support gun rights and the abolition of government, comments on reactions to two mass shootings in the U.S. from the day before.⁹

- (11) @_Kenziepuff: 'Conservatives posting tweets condemning mass shootings is more suspicious than not saying anything. No sane person "supports" those. Further proof mainstream conservatives are simply always reacting to the Left. Virtue signaling is more important to them than the American people.' (Twitter, 4 Aug 2019)

The user argues that it would be a waste of energy for the individuals in question to condemn the shootings simply to inform the public about their stance on the shootings, since it would be taken for granted that they were against the shootings. Hence, according to the user, there must have been some extra benefit to condemning the shootings: namely, appearing virtuous. Put in terms of (3d)'s *L*-alternatives factor, according to this user, the alternative of saying nothing, which would be less likely to look like an attempt to appear virtuous, still has much to offer, requiring less effort and being consistent with being against the attacks (which would supposedly be taken for granted). Therefore, the argument goes, there is good reason to think that individuals in question were

motivated by a desire to display virtuousness. This argument is not unassailable, however, as users responding to the tweet indicate. Multiple critiques argue that there are reasons for explicitly condemning the shootings other than attempting to appear virtuous, even if one's condemnation may be taken for granted. User @phoebesux6, for instance, writes, 'So expressing any feelings about people being dead is?? Virtue signaling?'

4.4. *Summary*

While it's rare for a speaker not to intend to suggest anything at all by the semantics of their utterance, what exactly the speaker does intend to suggest is not always clear, and can be both contentious and consequential. And the same principles that govern intention-attribution with respect to sociophonetic meaning, which center on beliefs about the agent's preferences, knowledge state, and appraisal of alternatives, apply just as well here. The preceding discussion thus echoes Franke et al.'s (2012) call for increased attention in pragmatic research to scenarios where the usual assumptions of Gricean cooperativity and overt intention aren't taken for granted.

Continuing along these lines, in the next and penultimate section of this paper, I turn to the role of intention in Burnett's (2017, 2019) SOCIAL MEANING GAMES models (SMGs) and the RATIONAL SPEECH ACT models (RSAs) (e.g Goodman & Frank 2016) to which they are related. These models have greatly enriched our understanding of how messages are conveyed in linguistic exchanges. At the same time, I will explain why the foregoing discussion points to a need to make the models more complex. For example, as discussed in Section 3 and contrary to what SMGs and RSAs are designed to account for at present, speakers sometimes hope their intentions will go unrecognized, certain social meanings depend on being performed in a way that seems natural, and hearers may conclude things from utterances that they think the speaker had no intention of suggesting.

5. INTENTION AND GAME-THEORETIC THEORIES OF LANGUAGE USE

5.1. *Background*

Central to game-theoretic models of language use and interpretation (e.g. Franke 2009, Frank & Goodman 2012, Burnett 2017) is the idea that speakers and hearers reason about each other's beliefs in choosing and interpreting utterances, and they know this about each other. A speaker *S* wishing to convey information to a hearer *H* attempts to select the utterance that appears to offer the best mix of being inexpensive and likely to convey exactly the information *S* wishes to convey. Similarly, observing some utterance *u*, *H* interprets *u* relative to *H*'s prior beliefs and the assumption that *S* selected *u* as just described. The interaction can be construed as a game in that there are agents employing strategy to achieve some goals. Influential in such game-theoretic research are RSA models, which situate language use and interpretation in broader theories of cognition as probabilistic reasoning (Goodman & Frank 2016). RSA models have been applied to a wide-ranging phenomena, offering means for testing quantitative predictions about language use and interpretation and delivering general insights into how these processes work (see Goodman & Frank 2016 for multiple examples).

Building on the insights and architecture of RSAs (as well as Franke 2009), Burnett (2017, 2019) broadened the scope of game-theoretic research on meaning to include social meaning via her SMG models. In SMGs, a speaker *S*, choosing which variant of a linguistic variable to use, tries to select the best variant for projecting the persona *S* wishes to project to the hearer *H*. *H*, in turn, tries to figure out what *S* is trying to convey—both parties taking into account *H*'s prior beliefs about *S*'s persona. The principal difference between RSAs and SMGs is that whereas RSAs typically involve signals/inferences concerning some question under discussion (Roberts 1996) based on an utterance's semantic content, SMGs involve signaling/infering something about *S*'s desired persona based on the indexical field (Eckert 2008) of the variant *S* uses.

5.2. *Cooperativity, reliable information, and presumption of intention*

Among the shared assumptions for RSAs and SMGs to date is a version of Gricean (1975) cooperativity. In these models, it is common ground among *S* and *H* that *S*'s goal in issuing a particular utterance is to provide as much relevant and reliable information to *H* as possible, *modulo* considerations of cost, and that *H*'s goal is to infer the information *S* is attempting to convey and update their beliefs with that information.¹⁰ In these models, then, both parties' interests are in effect the same as regards the utterance under consideration. As Burnett (2019: 11) puts it: '[B]oth players win if [*H*] correctly interprets [*S*]'s message, updating their beliefs accordingly, and they both lose if [...*H*] comes to believe something different about the world than that which [*S*] intended'.

Two points here merit attention. The first point is simply that these models are designed to capture only meanings that *H* believes *S* intended to signal. The second concerns the assumption that *S* will only attempt to convey reliable information. In SMGs, where personae and indexical fields are modeled as sets of properties, this assumption is operationalized by assuming that *H* rules out any potential persona for *S* that shares no properties with the indexical field of the variant *S* uses. For instance, modeling the indexical fields of *-ing* and *-in*' as {competent, delicate} and {incompetent, casual}, respectively, Burnett (2017: 258) notes that in SMGs, if *H* hears '*-ing*, they discard the possibility that the speaker is [both] incompetent and casual'.

The formal nature of SMGs has sharpened our understanding of the dynamics of social meaning by requiring explicitness about how things are supposed to work and the assumptions upon which the system rests. As may already be apparent, however, SMGs to date don't fully square with empirical observations about social meaning in other literature. In the next subsection, I will point out some of the challenges SMGs face, drawing on the preceding discussion. I will then show that RSAs to date face related challenges, made manifest in the light of SMGs' challenges, lessons from sociophonetic research, and the discussion herein.

5.3. Challenges for SMGs

A key set of difficulties for SMGs stems from the assumptions that *S* and *H* mutually hope *H* will correctly apprehend *S*'s intended message and update their beliefs accordingly, and that *H* assumes *S*'s selected variants always provide reliable information about *S*'s personae. It is certainly not always in *H*'s best interest to believe whatever *S* intends for *H* to believe. Rather, *H* would more generally be well served by trying to figure out what *S* is actually like. Whether *S* is truly competent or friendly or just attempting to appear so, for instance, can have important consequences for *H*.

Indeed, the discussion in Section 3.2.1 suggests that, *ceteris paribus*, differences in the nature of sociophonetic meaning and semantic content provide a reason for hearers to be quicker to doubt the reliability of an apparently intended sociophonetic meaning than the reliability of an utterance's semantic content. Franke et al. (2012) and McCready (2015) point out that if a speaker is found to willingly provide unreliable information, it can cost that speaker in terms of social capital (as in the case of the boy who cried, 'Wolf!'). Even amoral speakers thus have some incentive to tell the truth in the long run, particularly regarding easily verifiable matters of fact (though other considerations may trump that incentive). However, as also discussed in Section 3.2.1, given the performative and (inter)subjective nature of sociophonetic social meaning and its relative amenability to denials of intention by the speaker, it is relatively difficult to establish conclusively that a speaker suggested unreliable information via the phonetics of their utterances, not to mention that they did so deliberately. Therefore, we might generally expect the extrinsic incentive for speakers to provide reliable information to be weaker with sociophonetic meaning than with, say, the semantics of objective assertions—creating room for doubt about the reliability of what a speaker's phonetics might indicate about the speaker.

Moreover, predictions aside, the empirical research discussed in Section 3 has clearly shown that hearers don't always interpret a speaker's phonetics in the way they think the speaker intended.¹¹ Recall Campbell-Kibler's (2008) case of

‘Valerie’: ‘[Participants] who think she intends “intelligence” by her social cues react by seeing her as less intelligent’ (655). SMGs in their current form have no way of accounting for such cases. Given the assumption of cooperativity as implemented in SMGs, *H* assumes that *S* will provide reliable information, and if *H* thinks that *S* intends to signal some property via *S*’s chosen variants, *H* will ascribe that property to *S*.

Podesva et al.’s (2015) study of perceptions of U.S. politicians poses similar problems for SMGs. For instance, the authors found that Barack Obama’s speech was ‘rated as sounding more intelligent’ when his word-final /t/s were unreleased than when they were released. Research on word-final /t/ in English suggests that it is the released variant that would be more likely to be associated with intelligence, given its well-documented associations with traits like articulateness and learnedness (e.g. Benor 2001, Bucholtz 2001). But, again, given the assumption of SMGs that hearers will take speakers at their word, so to speak, we should be surprised to find a speaker being perceived as having a lower degree of some property (like intelligence) when they use a variant more closely associated with that property (like released /t/) than when they use another.

More generally, as discussed extensively in Section 3, that a hearer *H* believes that a speaker *S* intends to convey *c* via the phonetics of their utterance is neither necessary nor sufficient for *H* coming to believe *c* on the basis of *S*’s phonetics. *H* may infer things that *S* had no intention of suggesting; *H* may reject some meaning *H* thinks *S* intended to convey; and so on. Nor are such cases marginal in the realm of social meaning. But given SMGs’ assumptions about intentionality and cooperativity, they fall outside of SMGs’ scope in their present form.

5.4. *Summarizing, and looking ahead*

That said, though such cases of sociophonetic meaning aren’t accounted for by SMGs at present, they are presumably governed by many of the same general principles of probabilistic reasoning, social recursion, and strategic use and interpretation of linguistic resources that SMGs currently employ. It’s just that

such cases don't involve such strong assumptions about speakers and hearers' intentionality, common ground, and shared objectives. Rather, such cases call for a more complex model in which, with respect to a particular sociophonetic variant v , S might well be uncertain about, for instance: (i) whether H would interpret v as suggesting something about who S is; and (ii) whether H would interpret v as being motivated by a desire to suggest something about who S is. On the flipside, H may well be uncertain about: (i) whether S intended to suggest anything via v ; and (ii) how well the indexical field of v squares with S 's actual personality, stances, etc. Adding complexity along these or similar lines is necessary to cover the complex reality of social meaning. Franke et al.'s (2012) work on game-theoretic pragmatics without the strong assumption of cooperativity could prove rather helpful in this regard (see also Franke 2013). While that work focuses on inferences based in semantics, it proceeds from the fact that speakers and hearers need not have the same interests but may still apply strategy and reason in using and interpreting language. It is my hope that the principles developed herein concerning how speakers navigate uncertainty about intention will likewise help shape models of social meaning with broader scope.

Before concluding, I turn to a couple of related challenges for RSAs to date—challenges brought into relief by the foregoing discussion and insights from sociophonetic research.

5.5. *Related challenges for RSA models*

Campbell-Kibler (2008: 654) rightly points out that social meaning in situated use doesn't simply 'reside in the speaker's intention', thus distinguishing it from the type of meaning with which Grice (1957) was fundamentally concerned. But it is worth further pointing out that, in practice, meaning in context NEVER resides entirely in the speaker's intention, whether it's based in semantics, indexicality, or something else. Franke et al. (2012), for instance, examine a range of a cases in which a rational hearer might infer something from an utterance's semantics that the speaker didn't intend at all.

Indeed, lessons from sociophonetic research about gaps between intention and interpretation point to analogues within semantically based meaning that present challenges for models that assume cooperativity. Take for instance the phenomenon of ‘hypercorrection’ (Labov 2006 [1966]), whereby a speaker, apparently attempting to signal prestige, uses a prestigious variant of a variable more than would members of a more prestigious group with which the variant is associated. By overdoing it, speakers may end up being perceived as NOT having the relevant property (in this case, prestige). As noted above, SMGs cannot at present account for such cases; they assume hearers will regard speakers’ sociolinguistic performances as ‘truthful’ regardless of how convincing they really are.

This same dynamic can occur with semantically based meaning. Just as one can overuse a sociophonetic variant, one can assert something so many times or with such force as to lead the hearer to suspect that the assertion is false. We see this in Queen Gertrude’s reaction to the play within the play in *Hamlet*, in which her counterpart ‘protest[s] too much’ (III, ii): the play’s queen’s vows of faithfulness to her husband are so overdone as to seem insincere. Like SMGs, RSAs are not presently designed for cases in which a speaker intends to communicate *p* repeatedly, and the hearer recognizes as much, but, eventually, on the basis of the speaker’s repeated assertions, concludes *not p*.

More generally, RSAs aren’t presently designed for cases where the hearer believes that the speaker intends to communicate *p* but concludes not *p*. Kao & Goodman (2015) provide an insightful account of irony, which includes hearers concluding that the opposite of the semantic content of the speaker’s utterance is true. But in the cases their models account for, the hearer assumes that speaker intended for the hearer to come to that conclusion—that is, that the speaker wanted the hearer to conclude *not p* even though the semantics of their utterance was *p*—so the assumption of cooperativity remains.

Just as ‘protesting too much’ can raise suspicions about one’s sincerity, even a single instance of an utterance or variant can be delivered in such a way

that hearers take the opposite of an apparently intended meaning to be true. I have already discussed sociophonetic examples of this. But achieving one's goals via the semantics of one's utterance often requires a credible performance, too. Indeed—and this is crucial—all speaking is a performance of sorts, and using aspects of our utterances to change others' belief states in the way we intend requires getting our performances right.

Discourse surrounding a May 2019 *New York Times* report that, from 1985 to 1994, then private citizen Donald Trump lost more than a billion dollars (Buettner & Craig 2019: n.p.) provides an instructive case. When the hosts of the *Fox News* morning show *Fox & Friends* discussed the story, they depicted it in a generally positive light. Co-host Ainsely Earhardt said of the story:¹²

- (12) If anything you read this and you're like, wow. It's pretty impressive all the things that he's done in his life. It's beyond what most of us could ever achieve.

Many observers questioned whether Earhardt could sincerely believe that losing so much money is impressive (in a positive sense), and some suggested her delivery betrayed that she didn't actually believe her words. Aspects of her delivery drawing commentary included the focus of her gaze and the fact that she raised her hand in front of her mouth on the phrase *most of us could ever achieve*, taking an arrhythmic pause before the word *achieve*. Example (13) presents a handful of reactions from both media personalities and casual observers.

- (13) (a) Colin Jost of *Saturday Night Live's Weekend Update*: 'Come on, blond lady, even you— even you don't believe that. I mean you— you said the last part into your hand.'¹³
- (b) Twitter user @countdown2march: 'Lol they're not even trying anymore. Look at her face when she's saying that bullshit' (Twitter, 8 May 2019)
- (c) Late night TV host Jimmy Kimmel on *Jimmy Kimmel Live*: 'She almost couldn't say [achieve].'¹⁴

Each reaction suggests that something about Earhardt's delivery indicates that she doesn't truly believe her words. But the comments make no suggestion that Earhardt wants to come across that way. Rather, the interpretation seems to be that she intends to convey that she does believe what she's saying, but fails to do so. This is in contrast to cases of (cooperative) irony, where the speaker intends to suggest the opposite of what they say and wants the hearer to understand that.

Thus, just as appealing to the indexicality of a phonetic variant doesn't guarantee that one's intended effect will be achieved, so it goes with the semantics of one's utterances. Successfully making an explicit public commitment to a belief or preference (Condoravdi & Lauer 2012) means selling it, and sometimes further requires selling it without looking like you're trying to do so. Even expressives, which Potts (2007: 167) argues 'do not offer content so much as inflict it', can be unconvincingly performed. While it's true that the content of things like swear words isn't open to the same mechanisms of denial as, say, at-issue content, a hearer will only believe a speaker is in the emotional state conventionally associated with the relevant word if the delivery of that word is right. Again, while we may generally default to taking speakers at their word (and as Franke et al. 2012 point out, for good reasons), carrying off one's intended communicative goal means delivering a credibly sincere performance. This may especially be the case for expressions of emotion, opinion, or evaluation like (12), where, as with the case of social meaning (see Section 3.2.1), much of what's suggested cannot be independently verified by objective means.

In brief, though mismatches between what's inferred and what's intended may be more frequent or readily apparent with social meaning, semantics-based meaning is subject to the same dynamics. Developing a game-theoretic model of language use that can accurately predict inferences in cooperative contexts and also detect when utterances intended to be viewed as sincere are likely to be considered insincere is a tall order (see e.g. Franke et al. 2012 and Kao & Goodman 2015 for some steps in that direction). But humans' use and interpretation of language demonstrably involves both, among other complicating

factors. Not only do we sometimes doubt the truth of an utterance's semantics, but, relatedly, the cliché that it's not just what you say but how you say it is true, and these are facts that we as speakers and hearers must contend with every day. Ultimately, so, too, must our theories of language use and interpretation.

6. CONCLUSION

The overarching goal herein has been to clarify the nature of intention-attribution in general and what it in turn tells us about meaning. Section 2 laid the foundation with two key components. First was the principle of intention-attribution in (3), which, roughly speaking, states that our believing that an agent *A* intended to effect consequence *c* via action α is directly related to these four factors: (3a), believing that *A* would view *c* as desirable; (3b), believing that *A* would expect α to effect *c*; and (3c) and (3d), believing that any accessible alternatives apparently more (less) likely to effect *c* would be relatively unattractive (attractive) to *A* for independent reasons. Second was the observation, rooted in Grice (1957), that there are multiple ways for an action to suggest something to an observer—that is, to be MEANINGFUL: sometimes the relevant meaning is overtly or covertly intended, and sometimes it's not intended at all.

In the spirit of Campbell-Kibler (2008), Section 3 demonstrated that cases where it's unclear what, if anything, speakers intend to suggest sociophonetically are by no means rare. I showed that (3d)'s *L*-alternatives factor helps explain why sociophonetic meaning is especially amenable to being perceived as unintended: briefly, because speaking requires phonetics, there's nearly always the chance that the utterance's phonetic character was strictly in service of attempting to efficiently articulate the relevant expressions. I also showed that (3) illuminates social meaning more generally. For instance, in addition to according with previous sociophonetic perception studies, it explains why knowledge of a speaker's phonetic tendencies plays a pivotal role in determining whether a particular social meaning was intended.

Section 4 showed that semantically based meaning is subject to the same

dynamics of intention-attribution and uncertainty concerning intention as socio-phonetic meaning. The utility of (3) was again demonstrated in Section 4, where I showed, for instance, that the factors of intention-attribution in (3) undergird the kinds of arguments people make in debates about speakers' intentions.

Finally, in Section 5, I showed that game-theoretic models that assume cooperativity of the relevant kind miss much of the meaning utterances can bear and the complexity and uncertainty language users must grapple with in practice. This is perhaps especially clear with SMGs to date because, again, indexically based meaning is particularly amenable to 'no-intention' interpretations and particularly dependent on performance. But upon close inspection, one finds that many of the challenges SMGs face show up with semantics-based RSA models, too—underscoring the profoundly contingent nature of all meaning in practice.

Careful consideration of the role of intention in meaning foregrounds just how complex meaning-making is, and just how much uncertainty language users contend with. I hope that the present work has made this clear, and that the principles and analyses developed herein will prove useful to others in enriching our theories of meaning going forward.

REFERENCES

- Acton, Eric K. forthcoming. Pragmatics and the third wave: The social meaning of definites. In *Social meaning and linguistic variation: Theorizing the third wave*, Cambridge: Cambridge University Press.
- Ariel, Mira. 2004. Most. *Language* 80(4), 658–706.
- Austin, J. L. 1962. *How to do things with words* Second. Cambridge, MA: Harvard University Press.
- Bach, Kent. 2006. Review of Christopher Potts, ‘the logic of conventional implicatures. *Journal of Linguistics* 42(2), 490–495.
- Bach, Kent. 2012. Saying, meaning, and implicating. In *The Cambridge handbook of pragmatics*, 47–68. Cambridge: Cambridge University Press.
- Bargh, John A., Michelle Green & Gráinne Fitzsimons. 2008. The selfish goal: Unintended consequences of intended goal pursuits. *Social Cognition* 26(5), 534–554.
- Beauchamp, Zack. 2019. The Ilhan Omar anti-Semitism controversy, explained. *Vox* <https://www.vox.com/policy-and-politics/2019/3/6/18251639/ilhan-omar-israel-anti-semitism-jews>. 6 Mar 2019. Last accessed 30 Jul 2019.
- Benor, Sarah Bunin. 2001. Sounding learned: The gendered use of /t/ in orthodox Jewish English. In Tara Sanchez & Daniel Ezra Johnson (eds.), *U. Penn working papers in linguistics 7.3: Selected papers from NWA V 29*, 1–16.
- Blakemore, Diane. 2013. The expressive meaning of racial epithets: Towards a non-unitary account of expressive meaning. University of Brighton College of Arts and Humanities – Lectures in Language and Linguistics.
- Bucholtz, Mary. 2001. The whiteness of nerds: Superstandard English and racial markedness. *Journal of Linguistic Anthropology* 11, 84–100.
- Buettner, Russ & Susanne Craig. 2019. Decade in the red: Trump tax figures show over \$1 billion in business losses. *New York Times* <https://www.nytimes.com/interactive/2019/05/07/us/politics/donald-trump-taxes.html>.
- Burnett, Heather. 2017. Sociolinguistic interaction and identity construction: The

- view from game-theoretic pragmatics. *Journal of Sociolinguistics* 21(2), 238–71.
- Burnett, Heather. 2019. Signaling games, sociolinguistic variation and the construction of style. *Linguistics and Philosophy* 42, 419–450.
- Butler, Judith. 1993. *Bodies that matter: On the discursive limits of 'sex'*. Abingdon: Routledge.
- Campbell-Kibler, Kathryn. 2008. I'll be the judge of that: Diversity in social perceptions of (ING). *Language in Society* 37, 637–659.
- Campbell-Kibler, Kathryn. 2009. The nature of sociolinguistic perception. *Language Variation and Change* 21, 135–156.
- Clark, Herbert H. & Jean Fox Tree. 2002. Using 'uh' and 'um' in spontaneous speaking. *Cognition* 84(1), 73–111.
- Condoravdi, Cleo & Sven Lauer. 2012. Imperatives: Meaning and illocutionary force. In *Empirical issues in syntax and semantics* 9, 37–58.
- Du Bois, John W. 1993. Meaning without intention: Lessons from divination. In Jane H. Hill & Judith T. Irvine (eds.), *Responsibility and evidence in oral discourse*, 48–71. Cambridge: Cambridge University Press.
- Eckert, Penelope. 2008. Variation and the indexical field. *Journal of Sociolinguistics* 12(4), 453–476.
- Eckert, Penelope. 2012. Three waves of variation study: The emergence of meaning in the study of variation. *Annual Review of Anthropology* 41, 87–100.
- Eckert, Penelope. 2019. The limits of meaning: Social indexicality, variation, and the cline of interiority. *Language* 95(4), 751–776.
- Eckert, Penelope & Sally McConnell-Ginet. 2013. *Language and gender*. Cambridge: Cambridge University Press 2nd edn.
- Frank, Michael C. & Noah D. Goodman. 2012. Predicting pragmatic reasoning in language games. *Science* 336(6084), 998.
- Franke, Michael. 2009. *Signal to act: Game theory in pragmatics* ILLC Dissertation Series. Institute for Logic, Language and Computation, University of Amsterdam.

- Franke, Michael. 2013. Game theoretic pragmatics. *Philosophy Compass* 8(3), 269–284.
- Franke, Michael, Tikitu de Jager & Robert van Rooij. 2012. Relevance in cooperation and conflict. *Journal of Logic and Computation* 22(1), 23–54.
- Gaenzle, Martin. 2016. Meaning, intention, and responsibility Rai divinatory discourse. *Oral Tradition* 30(2).
- Goodman, Noah D. & Michael C. Frank. 2016. Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences* 20(11), 818–829.
- Grice, H. Paul. 1957. Meaning. *The Philosophical Review* 66(3), 377–388.
- Grice, H. Paul. 1975. Logic and conversation. In Peter Cole & Jerry Morgan (eds.), *Syntax and semantics*, vol. 3: Speech Acts, 43–58. New York: Academic Press.
- Horn, Laurence R. 1984. Toward a new taxonomy for pragmatic inference: Q-based and R-based implicature. In Deborah Schiffrin (ed.), *Meaning, form, and use in context: Linguistic applications*, 11–42. Washington: Georgetown University Press.
- Horn, Laurence R. 2004. Implicature. In Laurence R. Horn & Gregory Ward (eds.), *Handbook of pragmatics*, 3–28. Oxford: Blackwell Publishing Ltd.
- Horn, Laurence R. 2006. The border wars. In Klaus von Heusinger & Ken P. Turner (eds.), *Where semantics meets pragmatics*, Oxford: Elsevier.
- Kao, Justine T. & Noah D. Goodman. 2015. Let's talk (ironically) about the weather: Modeling verbal irony. *Proceedings of the 37th Annual Conference of the Cognitive Science Society*.
- Kao, Justine T., Jean Y. Wu, Leon Bergen & Noah D. Goodman. 2014. Nonliteral understanding of number words. *Proceedings of the National Academy of Sciences of the United States of America* 111(33), 12002–12007.
- King, Brian W. 2018. Hip hop headz in sex ed: Gender, agency, and styling in New Zealand. *Language in Society* 47, 487–512.
- Labov, William. 2006 [1966]. *The social stratification of English in New York City*. Cambridge: Cambridge University Press 2nd edn.

- Labov, William. 2012. *Dialect diversity in America: The politics of language change*. Charlottesville: University of Virginia Press.
- Lavandera, Beatriz R. 1978. Where does the sociolinguistic variable stop? *Language in Society* 7(2), 171–182.
- Levinson, Stephen C. 1988. Putting linguistics on proper footing: Explorations in Goffman's concepts of participation. In Paul Drew & A. Wootton (eds.), *Erving Goffman: An interdisciplinary appreciation*, 161–227. Oxford: Polity Press.
- Levon, Erez. 2014. Categories, stereotypes, and the linguistic perception of sexuality. *Language in Society* 43, 539–566.
- McCready, E. 2015. *Reliability in pragmatics* Oxford Studies in Semantics and Pragmatics. Oxford: Oxford University Press.
- Milne, A. A. 2009 [1926]. *Winnie-the-Pooh*. New York: Dutton Children's Books.
- Montague, Richard. 1970. English as a formal language. In Bruno Visentini et al. (ed.), *Linguaggi nella società e nella tecnica*, 189–224. Milan: Edizioni di Comunità. Reprinted in Montague (1974), 188–221.
- Montague, Richard. 1974. *Formal philosophy: Selected papers of Richard Montague*. New Haven, CT: Yale University Press. Edited and with an introduction by Richmond H. Thomason.
- Podesva, Robert J. 2011. Salience and the social meaning of declarative contours: Three case studies of gay professionals. *Journal of English Linguistics* 39(3), 233–264.
- Podesva, Robert J., Jermy Reynolds, Patrick Callier & Jessica Baptiste. 2015. Constraints on the social meaning of released /t/: A production and perception study of U.S. politicians. *Language Variation and Change* 59–87.
- Potts, Christopher. 2007. The expressive dimension. *Theoretical Linguistics* 33(2), 165–197.
- Potts, Christopher. 2015. Presupposition and implicature. In Shalom Lappin & Chris Fox (eds.), *The handbook of contemporary semantic theory*, 168–202. Oxford: Wiley-Blackwell 2nd edn.
- Roberts, Craige. 1996. Information structure: Towards an integrated formal theory of pragmatics. In Jae Hak Yoon & Andreas Kathol (eds.), *OSU working*

- papers in linguistics*, vol. 49: Papers in Semantics, 91–136. Columbus, OH: The Ohio State University Department of Linguistics. Revised 1998.
- Romaine, Suzanne. 1984. On the problem of syntactic variation and pragmatic meaning in sociolinguistic theory. *Folia Linguistica* 18, 409–438.
- Shakespeare, William. 2009. *The tragedy of Hamlet, Prince of Denmark* 1–138. Cambridge Library Collection - Literary Studies Cambridge University Press. doi:10.1017/CBO9780511704130.005.
- Smith, E. Allyn, Kathleen Currie Hall & Benjamin Munson. 2010. Bringing semantics to sociophonetics: Social variables and secondary entailments. *Laboratory Phonology* 1, 121–155.
- Sperber, Dan & Deirdre Wilson. 2004. Relevance theory. In Laurence R. Horn & Gregory Ward (eds.), *Handbook of pragmatics*, 607–632. Oxford: Blackwell.
- Stalnaker, Robert C. 2002. Common ground. *Linguistics and Philosophy* 25(5–6), 701–721.
- Stephens, Bret. 2019. Ilhan Omar knows exactly what she is doing. *New York Times* 7 March 2019.
- Tamminga, Meredith. 2017. Matched guise effects can be robust to speech style. *The Journal of the Acoustical Society of America* 142, 18–23.
- Weiss, Bari. 2019. Ilhan Omar and the myth of Jewish hypnosis. *New York Times* <https://www.nytimes.com/2019/01/21/opinion/ilhan-omar-israel-jews.html>. 21 Jan 2019. Last accessed 30 Jul 2019.
- Wharton, Tim. 2016. That bloody so-and-so has retired: Expressives revisited. *Lingua* 175–176, 20–35.

Author's address: Department of English Language & Literature
 Eastern Michigan University
 612 Pray-Harrod, Ypsilanti, MI 48197, USA
 eacton1@emich.edu

FOOTNOTES

1. Anonymized for submission.
2. Another way of formulating (4c), then, would be to say: ‘*S* intended *u* to suggest *p* to *H* by means of *H*’s recognition of this intention.’ I find this formulation less perspicuous, however, and the distinction has no practical bearing on the discussion herein.
3. This use of *natural* is not meant to coincide with Grice’s notion of NATURAL MEANING, though there is a practical relationship between them.
4. This discussion of the relation between intention and (un)natural ways of speaking highlights a link between (3d) and Horn’s (1984) DIVISION OF PRAGMATIC LABOR, whereby marked forms take on marked meanings.
5. Of course, even in (6) we don’t know for certain that the respondent positively wanted to suggest ‘not great’. For instance, the respondent may have felt strictly constrained to tell the truth and believed that ‘satisfactory’ was simply the most complimentary answer available given that constraint. See Ariel (2004), Horn (2006), Franke et al. (2012) for related discussion of the word *most*.
6. The same goes for the indexicality of variants of morphosyntactic variables (insofar as they are aptly analyzed as such; Lavandera 1978, Romaine 1984). Using, say, negative concord in English may be intended as a statement in its own right (Eckert 2019), depending on the speaker’s repertoire. A separate issue, however, is the SEMANTICS of morphosyntactic objects, to be discussed in Section 4.
7. Labov (2012: 22) discusses the related case of a travel agency with a ‘sign spelling out *cruisin*’, adding, ‘We understand this as an advertisement that we will have a better time *cruisin*’ than we would *cruising*’.
8. Busboys and Poets. 2019 Feb 27. <https://www.facebook.com/busboysandpoets/videos/353129905294312/>. Last accessed 29 July 2019.
9. This user’s account has since been suspended.
10. Even in RSA accounts of cases where hearers may conclude that an utterance is literally false, hearers still assume that the utterance conveys true and helpful

information. For instance, while a hearer may conclude that ‘This cup of coffee cost me \$50’ is strictly false as regards the cost of the coffee, it will still be taken to say something true and relevant about the speaker’s feelings about the cost of the coffee. As Kao et al. (2014: 12002) put it, in these models, it is ‘possible for a literally false utterance to be optimal as long as it is informative along the target dimension’.

11. It is worth noting that Burnett (2017, 2019) claims that the SMG ‘framework does not assume that all or even most aspects of message/interpretation selection or utility calculation are conscious or intentional’ (2017: 259), but Burnett uses the term *INTENTIONAL* in that quote differently than I use the term herein. Whereas one might gloss my sense of the term as ‘goal-directed’, Burnett’s seems to require a degree of conscious awareness. As the discussion in Burnett (2017, 2019) makes clear, speakers are indeed presumed to be making goal-directed decisions in selecting their utterances in SMGs. For instance, Burnett adopts the assumption that speakers ‘are trying to make the choice that will have the best chance of accomplishing their goals’ (2017: 248).
12. A video clip of the statement can be found here: <https://twitter.com/revrrlewis/status/1126094966609076225>.
13. Weekend Update: Trump Lost Over \$1 Billion. Aired 11 May 2019. NBC. <https://www.nbc.com/saturday-night-live/video/weekend-update-trump-lost-over-1-billion/3954440>. Last accessed 21 Sep 2019.
14. Jimmy Kimmel Mocks ‘Fox & Friends’ Defense of Trump Taxes Losses: ‘Will Make You Throw Up’. 9 May 2019. The Wrap. <https://www.thewrap.com/jimmy-kimmel-mocks-fox-friends-defense-of-trump-taxes-losses-will-make-you-throw-up/>. Last accessed 21 Sep 2019.