

Authentic Speech and Insincerity

Elmar Unnsteinsson

PREPRINT. PLEASE CITE PUBLISHED VERSION:

The Journal of Philosophy (forthcoming, open access)

([philpapers link](#))

Abstract

Many theorists assume that a request is sincere if the speaker wants the addressee to perform the act requested. I argue that this assumption predicts an implausible mismatch between sincere assertions and sincere directives and needs to be revised. I present an alternative view, according to which directive utterances can only be sincere if they are self-directed. Other-directed directives, however, can be genuine or fake, depending on whether the speaker wants the addressee to perform the act in question. Finally, I argue that this new perspective opens the door to a satisfying theory of authentic expression, for both assertive and directive utterances. Authenticity consists in the combination of genuine and sincere speech, for example, in the case of assertion, when speakers assert something which they both believe (sincerity) and want the addressee to believe (genuineness).

1 Introduction

People often use language to talk to themselves—reminding, encouraging, chastising—as well as to talk to others. This very innocent observation can, if allowed,

^{*}(✉) elmar.geir@gmail.com

^{*}I would like to thank a few people for comments and questions about this paper or its earlier incarnations: Hrafn Ásgeirsson, Maria Baghramian, Finnur Dellsén, Matti Eklund, Marinus Ferreira, Dan Harris, Thomas Hodgson, Eliot Michaelson, Carl Montan, James Norton, Jessica Pepp, Andreas Stokke, and Rowland Stout. I also thank two anonymous reviewers for this journal. This work was supported by the Icelandic Research Fund (grant no. 206551).

ripple through the dogmas in current speech act theory, crushing many long-standing, cherished ideas. Or so I will argue. Specifically, it turns out that a very influential theory of insincere speech and insincere directives seems plausible only when self-addressed speech is disregarded. But I can be insincere to myself as well as to others. Moreover, paying attention to self-talk, in this context, points the way to a new perspective on the authenticity of speakers, understood as the combination of sincerity and genuineness.

I start in §2, by spelling out the standard account of directive insincerity. This is, roughly, the idea that speakers are insincere in directing someone to do something if that is not something they desire. Here I propose two distinctions which will be important later. First, I distinguish the attitude the speaker wants to produce in the addressee, by making the utterance, from the attitude the speaker has in the production of the utterance. Second, focusing on the attitude produced, this should be broken down into the attitude-part and the content-part. I propose that the attitude is one of intention and that the content is a prescription rather than a proposition; thus, if I tell you to jump, and I am genuine, what I want is to produce in you the intention to jump.

In §§3-5 I have two aims. First, employing the distinctions already introduced, I pinpoint the exact sense in which the standard account assumes an asymmetry between directives and assertives. Second, I argue that this asymmetry only appears plausible because theorists have failed to consider the case of self-addressed speech. I conclude that any theory, like the standard account, which implies this asymmetry is thereby cast in doubt.

Finally, in §§6-8, I develop an alternative theory in some detail. As mentioned, this is based on the insight that there is a distinction between sincerity and authenticity in speech acts. Taking assertion as the paradigm case, speakers are insincere in asserting *p* if they do not believe *p*. But, moreover, the speech act is fake or not genuine if the speaker asserts *p* but does not want or intend the addressee, thereby, to believe *p*. The assertion is authentic, however, when the speaker both believes *p* and wants the addressee to believe *p*. This distinction is readily applied to directive utterances and self-addressed speech.

In this final part, I also show how this view supports the idea that so-called bald-face lies are not really lies, but rather inauthentic assertions. And I point out how the resulting theory of directive competence predicts that higher-order communicative intentions are involved in the maturation of pragmatic and linguistic competence from a very early stage of development.

2 The Standard Account

The standard account of insincere speech is often traced to John Searle (1969) and its basic structure has been endorsed by many since. Assertion is normally considered the paradigm case.

INSINCERE ASSERTIVE

S is insincere in asserting *p* if and only if *S* does not believe *p*.

Many philosophers have argued that even this simple account needs to be revised. Most of these proposals have to do with the role of the mental state of belief. Some have argued, for example, that the belief must be conscious or that it must be higher-order. But otherwise the basic idea has retained its currency (Chan & Kahane 2011, Eriksson 2011, Moran 2005, Ridge 2006, Schwartz 2020, Stokke 2014, 2018). As is well known, however, Searle also proposed sincerity conditions for other types of speech act. And it is commonly assumed in the literature that sincerity ought to apply to more than only assertive utterances (Kaplan 2004, Stokke 2014).

Here I focus on directives. Typically, directive speech acts are understood as utterances whereby the speaker directs the addressee to perform some action (Boisvert & Ludwig 2006). The standard account of directive insincerity, from Searle (1969, 66), is the following.

INSINCERE DIRECTIVE

S is insincere in directing *H* to *v* if and only if *S* does not want *H* to *v*.

This certainly strikes a chord in common sense and roughly the same idea was proposed earlier by R.M. Hare (1952, 13). Hare would say that *S* is insincere if *S* doesn't *intend* *H* to *v*. For my purposes, it will not be necessary to distinguish between different pro-attitudes—desire, preference, intention—in formulating the standard account, as the argument will apply to every theory of this general kind. Thus, the speaker is *directively* insincere, on the standard view, if and only if they lack some favored pro-attitude toward *H*'s *v*-ing.

To take a simple illustration, assume that I say to you,

- (1) (You) go home now!

In uttering (1), I will normally have performed the directive speech act of requesting or commanding that you go home now. In English the second-person

pronoun is optional. But if I utter (1) without actually wanting or intending you to go home it seems like something has gone wrong. On the standard account, I have failed to satisfy the sincerity condition for directives; to be sincere I must have a specific type of attitude, namely the desire or intention that the addressee perform the action I have requested. This pro-attitude appears to be part of what the speaker communicates by performing the directive act (Condoravdi & Lauer 2012, 38). At least, the speaker can reasonably predict that, on the basis of hearing (1), the addressee can infer that the speaker probably *wants* the addressee to go home.

But how exactly could the speaker fail to have the relevant pro-attitude in uttering (1)? A number of theorists have argued, in similar fashion, that this is impossible and, so, directives cannot be insincere, or are automatically sincere (Condoravdi & Lauer 2012, 43, Heal 1977, Siebel 2002, 129). My argument against the standard account (in §3) does not depend on this kind of impossibility and so, for sake of argument, I will assume that directives can be issued by speakers who lack the relevant pro-attitude. On the positive theory developed here, these will be fake or non-genuine rather than insincere, unless they are *self-directed*, in which case they could turn out to be insincere as well (see §6).

As Andreas Stokke (2014) has argued, the Searlean standard account can be generalized by saying, roughly, that insincerity consists in the communication of an attitude which the speaker lacks. I think this is a fruitful idea and will take it on board here as a working hypothesis. For example, the generalization will cover declarative and non-declarative utterances alike. To assert *p* insincerely is to communicate the belief that *p* while lacking that belief. To direct *H* to *v* insincerely is to communicate some pro-attitude towards *H*'s *v*-ing while lacking that attitude. Of course, the notion of what the speaker communicates is very broad, as a single utterance might communicate a number of different attitudes, and Stokke proposes that we restrict ourselves to consciously held attitudes. My argument exploits this feature to show, in effect, that the standard account fails to respect a structural symmetry between assertives and directives, which can be motivated independently. Roughly, if we make sure to separate the different contents conveyed by a directive utterance, we will alight upon a contender in the determination of insincerity. And this contender turns out to unearth the commonalities between insincere assertion and insincere directives, which have remained hidden to theorists so far.

But which contents, more precisely, are communicated by a literal directive? In the terminology of Condoravdi & Lauer (2012, 38), there is a distinction between what the directive *expresses* and what it *conveys*. On their view, (1) would

express the content, roughly: *the addressee goes home now*. But it also conveys *that the speaker wants the addressee to go home now*. Condoravdi & Lauer follow the standard view in thinking that the latter would be relevant for insincerity, concluding however that the speaker's desire is always present and thus directives are automatically sincere (p. 43). But I think the nature of the content expressed deserves more careful attention in this debate. As we shall see, it is important that acknowledge the role of the kind of attitude the addressee is supposed to bear to the expressed content.

At least since G.E.M. Anscombe (1957, §32) it has been common to separate the expression of propositions from the expression of *prescriptions* (e.g., Vranas 2008). At a first pass, this can be understood as a distinction between two groups of *attitude* speakers can bear to one and the same representation. The first is the group of *co-attitudes*, including belief and knowledge, which are supposed to represent the world truly or accurately. Any co-attitude that *p* aims to co-exist or otherwise correlate with the truth of the representation. The second is the aforementioned group of *pro-attitudes*, including desire and intention, whose satisfaction consists in *bringing about* the truth of the representation. In Searle's influential terminology, beliefs would have a word-to-world direction of fit and intentions a world-to-word direction of fit (1979, 3–4). Take 'milk,' for example, which could either belong to an inventory or a shopping list. In the first case, let's say, the representation is true if and only if there is milk in the fridge. In the second, it guides someone's action to make it so. To simplify, we can suppose that in both cases 'milk' expresses the same representation—something like: *that there is milk in the fridge*—and that the difference consists in the fact that thinkers can stand in different attitude relations to that representation. I can either believe or intend to make it true that there is milk in the fridge. I will drop this simplification shortly, however. Finally, we could then say that writing 'milk' on a list is either intended to activate the *p*-belief or the *p*-intention in the addressee, depending on whether it is an inventory or a shopping list. Directives seem, thus, to involve the expression of prescriptions and pro-attitudes.

To see this more vividly, consider three interesting features of directive acts. First, when I utter (1) to direct you to go home, I will not have said anything true or false. This can be tested by noting the infelicity of responding with 'That's true/false.' Second, it follows that I cannot have asserted anything, at least not directly, if assertion is understood as a speech act whose expressed content can be believed to be true or false. Certainly, there are theories according to which (1) expresses a modalized proposition, something along the lines of: you *should* go home now (Kaufman 2012). But one way to understand such a view is to dis-

tinguish, again, between what is expressed and what is conveyed. Thus, even if the latter involves a modalized proposition, there will be some level of expression at which the content is robustly non-truth-conditional, as indicated by true/false tests. As Craig Roberts (2018, 331) notes, we should prefer theories that can capture the non-assertive and non-truth-conditional nature of (1), if at all possible. Third, the expressed content of a directive is, in some sense, essentially indexed or anchored in the addressee. The directive can only be ‘satisfied’ if the addressee performs some action compatible with what the speaker communicates. For example, (1) is not satisfied by some audience-member other than the addressee going home now.¹

To respect these differences, we require something more fine-grained than the distinction between the content expressed and the content conveyed by a directive. A more detailed account will be spelled out in §6 below, but the proposal is inspired by a particular Gricean approach to distinguishing assertives from directives (Grice 1968, 1969, Schiffer 1972, Bach & Harnish 1979, Harris 2021). To start, we need to distinguish the attitude *produced* from the attitude of *production*. The normal function of a directive utterance is to (attempt to) produce—induce, activate, generate, whatever—an *intention* in the addressee. This is the produced attitude. The attitude of production is, rather, the speaker’s *intention to produce* the produced attitude. What I have been calling the content expressed is roughly equivalent to the produced attitude, while the content conveyed is roughly equivalent to the attitude of production. I go along with Jane Heal (1977, 191), and others, in holding that the speaker’s aim is to produce the *attitude* conducive to compliance, rather than merely the act of complying. This simple point calls for the structure of attitudes of production embedding produced attitudes.

We can see, already, how this picture suggest a strong analogy between directives and assertives: the normal function of the latter is to produce *beliefs*—or other co-attitudes—in exactly the same sense. Although this is not the focus of the current paper, and my arguments can play with other pro-attitudes, I will tend to talk in terms of the production of *intention* in the case of directives.²

¹According to Zanuttini et al. (2012), the imperative clause is part of the more general ‘jussive’ clause type, which also includes promissives and exhortatives (these are all syntactic categories, not speech act theoretic). Promissives have the speaker as subject while exhortatives are addressed to both speaker and hearer as a group. My argument only concerns ‘literal directives,’ understood as directive speech acts where what the speaker utters is an imperative clause.

²Harris (2021, 1069–1070) argues specifically for this view, on the basis that otherwise we could not explain the infelicity of directing addressees to do something they lack control over. To take his example, I cannot felicitously command you to stop having nightmares, because you

Finally, happily, this helps to further analyze the cognitive effect produced by directives. Assume that '*I*' stands for the attitude of intention and ' π ' for the *prescriptions* embedded by the intention. On this picture, when I say (1) to you and mean what I say literally, what I *mean* or *express* is some mental state of the form $I(\pi)$. That is to say, I mean for you to form the intention to go home now (Heal 1977, 191). We can also say that π is what I mean (prescribe, express), but that's merely a less informative specification, since I must also mean some attitude relation rather than another. I don't mean that you are supposed to *fear* or *hope* something, for example (see, e.g., Grice 1991, 63–64). In the terminology adopted from Condoravdi & Lauer above, $I(\pi)$ is the form of the content expressed (or produced), while one part of what is *conveyed* will be the speaker's intention that the addressee form the $I(\pi)$ -attitude, this being the attitude of production. In the positive account developed in §6 below I will call the latter the speaker's *effective* intention.

This framework fits nicely with the three features articulated above. Intuitively, π is an action type (*to go home now*) and $I(\pi)$ is an intention to perform a token action of that type. In formal semantics we can, as is currently popular, think of the denotation of the imperative clause-type as a property indexed or applied to the addressee. For example, Nate Charlow (2014) thinks of the denotation as a property of a *plan*, namely the property the plan has just if it requires that you go home now. A plan is a set of intentions, obeying basic principles of practical rationality. If you adopt a plan with this particular property, which might be the upshot of someone saying (1) to you, that part of your plan is satisfied if and only if you go home now. Among other things, this would help to explain why (1) and 'Don't go home now!' seem logically inconsistent, without being either true or false. Roughly, to adopt a plan with both properties is inconsistent in virtue of violating basic requirements of rationality (see also Charlow 2018, Harris 2021, Portner 2007, 2018, Roberts 2018). The property denoted by the clause partly determines the action type intended by the speaker in issuing the directive utterance. Anyway, neither the action type nor the intention seem to have intuitive truth conditions and the formation of an $I(\pi)$ -state—e.g., the intention to go home now—is more like a decision to act than a judgment about something being true. And this is exactly what a directive utterance, on this view, is supposed to achieve; the addressee's decision to form the attitude the speaker intends to produce.

cannot form the intention to stop, at least not directly. Portner (2004), Condoravdi & Lauer (2012), and Starr (2020) would rather talk in terms of influencing (or expressing) desires or preferences.

3 Why the Asymmetry?

In this section I will argue that the standard account of directive insincerity introduces an implausible asymmetry between sincere directives and sincere assertives. Later, I will explain why the asymmetry is implausible but it arises, I believe, because theorists tend to consider only *other-directed* speech in their accounts. No surprise, as many have assumed that the primary function—in some sense—of language is interpersonal information exchange. Even so, speakers engage in episodes of self-addressed private speech all the time, for various communicative or non-communicative purposes (see, e.g., Gregory 2016, Langland-Hassan & Vicente 2018, Winsler, Fernyhough & Montero 2009, Unnsteinsson 2022a). The current argument does not require any specific assumptions about the nature or function of inner or private speech, only the obvious fact that it exists.

Following this line of thinking, let's call speech acts 'paradigm' if they are primarily directed at an addressee other than the speaker. Now, there is one very striking difference between paradigm assertives and paradigm directives. According to the working hypothesis adapted from Stokke, the former will be sincere in virtue of the speaker's own private relation to the produced attitude. That is to say, if I want to produce the belief that *p*, I am sincere if and only if I myself *believe* that *p*. But a parallel definition appears incoherent in the latter case, because the produced attitude is not a potential state of the speaker at all. What I mean or express in uttering (1) is, if our assumptions so far are right, an intention to go home now, and the intention is supposed to be formed by you, the addressee. We can rummage around in my mind all we want, but we'll never find there an intention to go home now which will be relevant to the determination of whether or not my utterance is made sincerely. For example, my utterance does not become sincere if *I* intend to go home now. And I cannot even form *your* intention to go home now, because your intention has the function of guiding and controlling *your* conduct in a manner in which my intentions cannot. Certainly, there is an action type both of us can intend to perform, namely to go home now. But our intentions in performing the same action type are still essentially different in that my intentions guide my actions and not yours, and vice versa.

As indicated before, this difference generalizes to any pro-attitude we might want to use in theorizing about directives. Take preferences for example, so that the attitude produced will be: *your preference to go home now*. Again, there is no attitude in my mind which matches or is identical to your preference to go home now, such that my attitude will plausibly determine the utterance's

sincerity conditions. I might prefer to go home now, but that's irrelevant. I might prefer that you go home now, but that's not the attitude produced (it's not even an attitude of your's). Of course, the option remains to say, in response, that sincerity is determined by the attitude of production and not by the attitude produced. But my response will have to wait until we get to the argument that we should expect structurally analogous speech acts to have structurally analogous properties (§5).

Anyway, it follows that the insincerity of a directive cannot be determined by a direct mismatch between the $I(\pi)$ -attitude produced and the producer's own $I(\pi)$ -attitude. That is to say, I am not insincere in virtue of lacking the pro-attitude I aim primarily to communicate and produce. Rather, on the standard account, I am insincere in virtue of lacking some second-order attitude which embeds the produced attitude; that I *want* you to form the $I(\pi)$ -attitude (this would be the attitude of production). This is the asymmetry between insincere assertives and insincere directives on the standard account. Because assertions are supposed to produce belief in propositions—or so we will assume here—rather than intentions to act, they are very different. My belief that p and your belief that p are both true if and only if p , so they are essentially the same even if realized by different people. That is to say, their satisfaction conditions, or truth conditions, are identical. What I assert is something I myself can either believe or fail to believe and, so, my own first-order relation to the proposition expressed may, plausibly, determine my insincerity.

To drive the point home, notice that the asymmetry survives a comparison between addressee-indexed assertions and directives. So, the asymmetry does not derive from the fact that directive content mandates indexing, while for assertives this is optional.

(2) You went home yesterday

By assumption, if I say (2) to you, I speak insincerely if and only if I don't believe that you went home yesterday. No one would be inclined to think that my insincerity in uttering (2) is due to the fact that I don't really want you to believe that you went home yesterday. This indicates that the apparent asymmetry does not arise because the content expressed is somehow indexed to the addressee.

Jane Heal (1977) and Mark Siebel (2002) argue that various other asymmetries between assertives and directives indicate that there really is no such thing as an insincere directive. Condoravdi & Lauer (2012) propose that sincerity is automatic for directives. By contrast, the standard account assigns blame to a co-conspirator. Since the real offender is not available, we get the next best thing,

namely, the speaker's second-order desire or intention to produce the intention to π in the addressee. This is the attitude of production, which always embeds the attitude produced. And, as already mentioned, this may seem like the right conclusion. If I order you to go home when I really want you not to go home, I seem—if this is possible—to be speaking insincerely or dishonestly. But why exactly should we accept this particular asymmetry? On reflection, as I shall now go on to argue, it seems unmotivated.

4 The Argument from Self-Directed Speech

Now, let's consider utterances that fail to adhere to the paradigm of other-directedness. Specifically, I argue, it turns out that *self-directed* directives can be sincere or insincere just like assertives and for the same reasons. If I plan to go home at 3pm and, when I see that it's now exactly 3pm, I say (1) to myself, I am being sincere to myself, and literally so. The content of what I say or mean matches my own intentional state, and I utter (1), plausibly, to cause the activation or further operation of that state, for example in working memory. To write 'milk' on the shopping list can also be used as an example here. If the list-writer is directive sincere, they in fact intend to buy milk and, presumably, want that intention to be re-activated at some later time. If the shopping list is exclusively addressed to someone else—here I ignore the interesting case of directives which are simultaneously self- and other-addressed, sometimes called 'exhortatives' (see footnote 1)—sincerity of this sort is impossible, for reasons already canvassed. On the other hand, if the writer has no intention at all to buy milk, neither at the time of writing nor at the expected time of reading, the utterance was literally insincere.

More precisely, if π is the directive *content*, the prescription expressed, and $I(\pi)$ is the attitude I want to produce by my utterance, there can be a perfect match between the attitude produced and my first-order mental state. I utter (1) to produce the intention to π , while simultaneously having an intention to π . The attitude of production is the more complex intention to produce the intention to π . To make this even more obvious, perhaps, imagine that I write (1) in a self-addressed email which is timed to appear in my inbox at 3pm. Even as an interpreter, I must recognize my own prior intention to produce an intention to go home, in myself. Here, sincerity and insincerity work in exactly the same manner as in paradigm, other-directed assertives.

Imagine, to take a different example, that someone appears at the edge of jumping out of an airplane with a parachute on, but really has no intention of

going through with it. It seems coherent to suppose that they utter ‘Jump!’ wanting thereby to produce an intention to jump in themselves, even if they simply do not intend to jump, and never intended to jump. For this to be rational, the speaker only needs to believe that it is *possible* that the utterance will be effective in producing the intention, not that it is all that likely. Perhaps the example is more plausible when the speaker knows that others can hear these self-directed commands—the display is also supposed to inform others of something—but that doesn’t affect the main point. The point is that directives can be insincere with respect to what the speaker says or expresses, rather than merely with respect to some second-order desire embedding the $I(\pi)$ -attitude meant by the speaker. Addiction, weakness of will, and wishful thinking can also provide familiar examples. A smoker could have a first-order desire to smoke and a higher-order desire not to have that desire. A self-directed utterance of ‘Stop!’ might thus be intended literally to activate the intention not to pick up the cigarette.

Why exactly is this a problem for the standard account? Well, it seems like the account only allows for two easy reactions, both dead-ends. First, one could say that there is no such thing as an insincere directive which is self-addressed or self-directed in the requisite sense. But this flies in the face of common sense and a great deal of empirical evidence. Inner or private speech is a well-studied phenomenon and directives are certainly thought to play a significant role, for example in sport psychology (Hardy 2006). If so, we should predict that insincerity may feature as part of that role, just as it does when we lie to ourselves in self-directed assertions.

Second, and more promisingly, proponents of the standard account could insist, still, that self-directed insincerity consists in the speaker’s lacking the relevant second-order desire. We would need an explanation of why this is not true also in the case of assertion, that is, why is my self-addressed assertion then not insincere in virtue of me not wanting myself to believe what I assert? Even so, the idea might be to question the very distinction between second-order and first-order states in self-directed speech acts. What exactly is the difference, one might ask, between S wanting S to (intend to) go home now and S wanting or intending to go home now?

To examine this distinction in self-directed insincerity we need an example where the second-order desire is present while the speaker lacks the first-order desire or intention. And this was in fact the structure of the skydiver’s case. There is no intention to jump but, still, there could be a desire to produce the intentional act of jumping by way of the self-directed command. The idea was to say that this is an insincere directive, because the *sincere* case would be one

where the skydiver really intends to jump all along and makes the utterance as a reminder or piece of self-encouragement. Thus there would be a perfect match between the speaker's attitude in making the utterance (*intention to jump*) and the attitude produced or communicated (again, *intention to jump*). The presence of the second-order desire—the attitude of production—does not make the utterance sincere, if the difference between the first- and second-order states is a real one. We can compare this to assertion for illustration. Even if I want myself to believe that I'm not bald, saying to myself in the mirror 'I'm not bald,' I'm still lying to myself if I really believe that I'm bald. I conclude that the standard account implies an asymmetry between insincere directives and insincere assertions. But some might argue that this is not a real problem.

5 Is the Asymmetry Implausible?

Why should we think that the asymmetry is implausible?³ It is certainly not incoherent. Basically, the default assumption should be that the cognitive mechanism of pragmatic competence is uniform across the capacity to produce assertions and directives. On the Gricean view, to be explained in more detail later, pragmatic competence is the capacity to perform speech acts with a set of specific audience-directed intentions. Three clause types are normally considered universal or near-universal in human languages; declarative, imperative, and interrogative (Sadock & Zwicky 1985, Roberts 2018). Data from language acquisition suggests that the first two are acquired at the same or similar stages of development (Aikhenvald 2010, 325–330). This supports the assumption that the capacity to perform the corresponding speech acts is explained by one and the same cognitive mechanism. Besides, positing a single structure for both capacities—insincere assertives and insincere directives—is more parsimonious.

On this view, we already posit the relevant mechanism to explain assertive competence, by which I mean that even very early in the development of pragmatic competence, speakers can perform normal assertive acts, and that this is due to some mechanism in virtue of which the relevant intention-types can be formed and employed in the control of planned action. My claim here is that the default assumption should be that directive competence recruits exactly the same resources. If so, we should expect the attitude of production to have a certain structure. The speaker has (1) a pro-attitude towards producing a particular (2)

³I would like to thank a reviewer for this journal for raising the question, which gave rise to this section.

pro-attitude in the addressee. If this is right, we should expect speech acts to be insincere in virtue of a relation between (2) and the speaker's own pro-attitudes. But this only seems possible in self-directed speech, and this point has not been noted in the literature to date. The current proposal has the virtue of identifying a parameter which immediately explains the asymmetry that remains, which is precisely the distinction between self and other.

One might wonder, still, how this could be so. Will we not still predict the same asymmetry for all sorts of other speech acts, like questions, promises, and apologies? If so, it would seem like we should not be too bothered by it in the case of directives.

In my view I could have replaced directives with questions and made many of the points I have made so far. If we accept recent arguments from Whitcomb (2010), Friedman (2013), and Carruthers (2018), we should think of interrogative attitudes as 'basic' like intentions and beliefs. Further, interrogative clauses seem to be a human linguistic universal specifically suited for the expression of this attitude. But there are two important differences. First, interrogative attitudes are more controversial than intentions. Second, interrogative attitudes are easily shared like assertive (i.e., belief-like) attitudes. If it is my intention to produce in you curiosity about whether *p*, by uttering an interrogative, I might be insincere in virtue of not being curious myself about whether *p*. Now, I am not denying that intentions can be shared in some sense, only that a shared intention is what is directly meant or expressed by uttering a literal directive.

Promises and apologies are different because they are more likely to be relative to factors like culture, convention, and history, which probably has effects on the cognitive mechanisms involved and the process of acquisition. But we do make promises and apologies to ourselves, which is certainly relevant and interesting. It is equally clear that apology, for example, is not a conventional device for the production of a distinctive attitude, different from intention or belief. To simplify, assume that apology is associated with the attitude of regret. Still, apology is not in the business of producing regret but, more plausibly, it produces the *belief* that the speaker feels regret. A simplified account might be as follows (cf. Bach and Harnish 1979, 51–2):

APOLOGY

S apologizes for *v*-ing by uttering *X* if *S* thereby intends the addressee to believe that *S* regrets having *v*-ed.

Note that this allows for the idea that insincere apologies are apologies. And if this is right, the Gricean theorist is free to posit additional structure into the

attitude embedded by the effective intention (the attitude of production), and this seems to be a very satisfying way to capture the difference between apology and directives. Most importantly, this fits the schema I have adapted from Stokke; speech acts are insincere in virtue of the speaker lacking the produced attitude. I take this as further reason to accept the schema as a general theory of insincerity across different types of speech act.

There are many reasons, however, not to focus on highly conventionalized speech acts like apology or promising, only one of which is how they can vary between groups of speakers. The most pressing reason is the added complexity in the accompanying notion of insincerity. Is the apology insincere in virtue of the speaker's lack of regret, lack of belief that they are regretful, or both? Does the belief, the regret, or both, need to be consciously accessible? The same questions will arise on a similarly Gricean story about promises (again, very much simplified from Bach & Harnish 1979, 50):

PROMISE

S promises to *v* by uttering *X* if *S* thereby intends the addressee to believe that uttering *X* makes it obligatory that *S* intend to *v*.

Although I think something *like* this must be right, I will not argue for it here. I only want to point out that the theory retains a clear distinction between the attitude of production (i.e., effective intention) and the produced attitude, and that this allows for a general explanation of insincerity which is grounded in a theory of more basic speech act types. The insincere promise may consist in *S*'s lacking the belief that the utterance puts *S* under any obligation. But it seems like the promise could also be insincere in virtue of *S* lacking the intention to *v*. That's not a big problem, because in that case lacking the intention might be part of the explanation why *S* lacks the belief in question. Note also that these accounts of apology and promises adhere to specific Grice-inspired methodological strictures.⁴ That is to say, they are *speaker-centric*; promises and apologies are insincere in virtue of something about the speaker—and what the speaker means—not in virtue of something about the addressee's rational or normal reaction. That being said, promises and apologies are sometimes embedded in highly institutionalized and conventionalized practices which can have significant consequences for our taxonomy.

⁴See, for example, Bach (2005), Fodor & Lepore (2004), Neale (2016). I will turn to this matter again in §7 below.

More to our purpose, it seems like the potential generality of this Gricean account adds greatly to its plausibility. That is to say, we seem justified in thinking that insincere speech, generally, consists in a mismatch between the produced attitude and the speaker's own attitudes. Since there is also good reason to suppose that literal directives are conventional devices for producing intentions, just like literal assertives are supposed to produce beliefs, this general structure appears to fail in the case of insincere directives. Further, if what I've argued in this section is right, we can conclude that this structural disanalogy or asymmetry, implied by the standard account, is implausible. We are justified in looking elsewhere. But what is the alternative?

6 Authenticity and Sincerity

What we need is a distinction between sincere speech and *authentic* speech. To start, distinguish two intentions speakers have in performing communicative speech acts, according to the broadly Gricean framework mentioned above. First, there is the *effective* intention, which is the intention to have some specific cognitive effect on the addressee. This is what we have, so far, been calling the attitude of production. It is roughly equivalent to the second-order desire in the standard account of directives, applied to all communicative acts. For example, if I tell *H* that *p* by uttering something I will normally intend *H* to form the belief that *p* or, at least, to form some specific co-attitude *A* with respect to the proposition that *p*. The cognitive effect intended by the speaker, then, is going to be some *A(p)*-state formed by the addressee, in response to perceiving the utterance.

Second, on this framework, speakers will also have a *signaling* intention, which is the intention to get the hearer to recognize the effective intention. As Scott-Phillips (2015) would put it, the speaker wants to signal the signalhood of the utterance. The addressee's recognition of the signaling intention is then considered sufficient for *communicative* success (e.g., Grice 1989, Neale 1992, Sperber & Wilson 1986/1995, Schiffer 1972, Simons 2017a). When addressees satisfy the speaker's effective intention—by actually forming the attitude the speaker intends to produce—the speaker is more than merely communicatively successful. Combinations of effective and signaling intentions can be called communicative intentions.⁵

⁵Here I am only slightly modifying the standard terminology in relevance theory (Sperber & Wilson 1986/1995, Wilson & Sperber 2012, Carston 2002), where 'effective' is 'informative' and 'signaling' is 'communicative.' I think 'effective' is better here because directives are not

To generalize this theory, my act of trying to communicate $A(p)$ will be sincere if and only if I bear the attitude A to p . The same schema will apply to directives—with the qualification about the difference between self and other, to be discussed below—by replacing the cognitive effect with $I(\pi)$. Furthermore, my $A(p)$ -assertion, directed at addressee H , is *genuine* with respect to $A(p)$ if and only if I really intend to produce the $A(p)$ -state in H . On this view, genuineness is determined by the effective intention and my communicative act is *fake*, with respect to $A(p)$, if the corresponding effective intention is lacking. Finally, we can say that the assertion is *authentic* if and only if it is both sincere and genuine. It is inauthentic, then, exactly when it is both insincere and fake, with respect to one and the same cognitive effect.

An utterance can be fake but sincere with respect to the same attitude, for example in sincere irony (cf. Stokke 2018, 174). Consider Grice's example of uttering,

(3) He's a fine friend,

ironically or sarcastically, to mean that someone is the opposite of a good friend (1989, 34). On his theory, the speaker merely *makes as if* to say that Harald, to give him a name, is a fine friend—so, doesn't really *say* this—and thereby conveys that something else is really meant. What the speaker means is often thought to be, in some intuitive sense, the opposite of what was made as if to say.⁶ Mandy Simons (2017b) would say, in this case, that p is what the speaker merely *expresses*—without saying or meaning—and Stephen Neale (2016, 258) describes the phenomenon in terms of *play-saying* and *play-meaning* that p . Let's assume, then, that what the ironical speaker play-asserts by uttering (3) is a cognitive effect of the form $A(p)$, namely the belief that the man is a fine friend. Now, either the speaker believes p or does not believe p . In standard cases of irony, speakers do not believe the proposition they play-assert; instead they only believe what they really mean (that he's a jerk). But this is not essential to irony on the Gricean view.

That is to say, I could utter (3) ironically to communicate my belief that Harald is an awful friend, while believing that he is a fine friend (the $A(p)$ -state). If so, my utterance is *fake* but *sincere* with respect to what is play-asserted, namely $A(p)$.

obviously a matter of information exchange.

⁶I will not argue for this theory of irony or sarcasm here, but see Elisabeth Camp (2012) for a helpful discussion. For other theories of irony see, for example, Currie (2006), Lepore & Stone (2015), Wilson & Sperber (2012, ch. 6.)

The utterance is fake in this manner because I lack the effective intention to produce, by my utterance, the $A(p)$ -state in the addressee. But it is sincere with respect to this very state because it matches my own attitude. Importantly, it doesn't follow that I have no effective intention at all, the effective intention just involves a different cognitive effect. If this is right, irony can be fake and sincere with respect to what the speaker makes as if to say. Actors who happen to believe what they say on stage could also be engaged in fake but sincere assertions, in virtue of lacking the corresponding effective intentions. When actor A says to B , on stage, 'I am your mother,' she doesn't want or intend to activate in B the belief that A is B 's mother. It's only part of the script. But if, by some historical accident, the belief is true and A is desperately—given the unlucky circumstances—trying to hide it from B , the utterance is sincere but fake with respect to what A play-asserts.

Normally, however, irony is inauthentic with respect to one and the same content, for the speaker neither believes nor intends to produce the belief that, in this example, Harald is a fine friend. Rather, the speaker intends the hearer to recognize that $A(p)$ is *not* meant, so some other cognitive effect can be inferred as what *is* meant (that Harald is a jerk).

These distinctions are readily applied to directives. Directives can be sincere or insincere only when self-directed, as already argued. But other-directed directives are *fake*, exactly like assertives, when the speaker lacks the relevant effective intention. This is the case, for example in directive irony. To borrow an example from Alexandra Aikhenvald (2010, 234), imagine that your partner comes home late one night and, in a drunken rampage, starts trashing the place. You say,

(4) That's right. Break everything in the house!

There is a cognitive effect here which you merely express or give voice to, namely your partner's intention to break everything in the house. If uttered ironically, (4) is really intended to produce a very different effect in your partner. The opposite effect, it may seem.⁷ Thus, the directive is fake with respect to

⁷Grice's (1975) account of implicature would require some slight modification to accommodate the notion that speakers can implicate something by a directive, if directive contents lack truth conditions. The *Maxim of Quality* is formulated in terms of the speaker's attention to *truth* and *evidence*. For one model of how to do this, see David Braun's (2011) account of how speakers can implicate by asking questions. Part of his proposal is to replace the *Maxim of Quality* with a more general *Maxim of Sincerity* (2011, 589).

this merely expressed or play-directed effect. For concreteness, let's assume that there are two cognitive effects at issue:

- (5) *I(break!)*
- (6) *I(don't break!)*

On this view, by uttering (4), your partner fakes (5) in order to convey a different cognitive effect, namely (6). As an act of communication, then, the utterance is successful if and only if your partner recognizes your effective intention, namely that they shouldn't wilfully break everything in the house (6).

Of course, this theory of irony is controversial, but serves as a convenient illustration (see footnote 6). The distinction between fake and genuine speech acts, understood in terms of effective intentions, applies much more widely. Consider, briefly, directives with foregone conclusions. If I'm tied to a chair while some hoodlums rummage through my stuff, I could say, 'Take the money! What do I care!'. The directive definitely lacks the associated effective intention and is thereby fake or non-genuine. That is to say, I faked a directive with the content, roughly: *intention to take the money*. Here the speaker's further communicative purpose is less transparent, but the utterance is fake rather than insincere with respect to this particular content. It is not insincere, as already argued, because the attitude is not available to me as a first-order mental state. It is important to notice that speech acts are defined, here, as fake, insincere, or inauthentic only with respect to some specified cognitive effect or other. Faked effects are in the service of the effects actually intended by the speaker.

So far, I have avoided the topic of *lying*, which has received a lot of attention recently (Carson 2010, Michaelson & Stokke 2018, Saul 2012, Stokke 2018). In the literature on lies, theorists have become interested in the notion of *bald-faced* lies, which seems relevant here. Lies are bald-faced, I take it, when the speaker does not disguise the fact the utterance is a lie. Many argue, on this basis, that lies do not require the speaker to have an intention to deceive the addressee. It is common ground that the speaker doesn't believe what is said, so deception doesn't seem to be part of the speaker's goal in making the utterance (Carson 2006, 2010, Fallis 2012, Saul 2012, Sorensen 2007, Stokke 2018).

My argument so far would suggest, instead, that bald-faced lies are assertions that are both fake and insincere. So, relative to a specific cognitive effect, $A(p)$, the speaker lacks the intention to produce that effect in the addressee and, at the same time, does not bear the attitude A to the proposition p . This agrees, more or less, with the views of a group of philosophers who argue that bald-faced lies are

not really assertions at all, essentially on the basis that the relevant effective intention, or something very much like it, is missing (Harris *forthcoming a*, Keiser 2016, Maitra 2018). On the theory proposed here, bald-faced lies are cases where the speaker's expression of a particular attitude is not authentic. So, for example, if a defendant takes the stand and swears to the jury that they did not commit a crime, they might be speaking inauthentically, in part because everyone already knows that they are guilty. Even so, the defendant might not want to be on the record admitting their guilt. But if the truth is already common ground, it seems like the speaker neither believes *p*, nor intends the audience to form the belief that *p* on the basis of the utterance. Thus, some other intention or purpose must be guiding the speech act.

Finally, we can now see clearly how this theory can unify the notions of sincerity and authenticity across speech act types. The differences only show up in highly predictable situations; there is no such thing as other-addressed directive insincerity, only non-genuineness or faking. Faking is understood, both for assertives and directives, in terms of the speaker lacking the effective intention with respect to a particular attitude. Directive sincerity still exists, but it is only defined for self-directed speech. I have not yet given any examples of directive speech acts that fail to be *authentic*. They are easy to come by, however. Consider the timid skydiver again, and call him Marshall. He says (7) to himself but, this time, he does so only in order to placate the others watching. He wants everyone to think that he is seriously trying to overcome his fear, but has long since decided that he'll sooner die than jump from an airplane.

(7) Jump!

To explain Marshall's overall plan, it seems that we must posit a specific cognitive effect which is merely faked or play-meant, something like (8).

(8) *the intention to jump*

On this picture, he utters (7) to himself and part of his goal is to fake (8) as the effect or attitude he thereby means to activate or produce in himself. Marshall's real purpose, however, is to make the others believe that he really intends to jump, or intends to incite himself to jump. These others are what Herbert Clark (1992, 202) would call 'side participants' in the conversation. Marshall is both the speaker and an addressee, but (7) is uttered in part to inform side participants that he has performed a speech act intended to produce the attitude in (8) in himself. Of course, the side participants are not supposed to recognize that

(8) is the *faked* effect of uttering (7). Indeed, Marshall knowingly lacks the effective intention of producing (8) in himself by uttering (7) and does not possess the attitude involved. Thus, the directive act is inauthentic with respect to (8). Plausibly, the utterance is genuine but insincere with respect to some cognitive effect Marshall is trying to generate in the minds of side participants. Here's a possible candidate:

(9) *the belief that Marshall intends to jump*

That is to say, Marshall may utter 'Jump!' to produce (9) in the others, while (9) is not one of his actual beliefs. If so, the utterance is insincere but not fake with respect to (9).

This theory of authentic speech constitutes a serious and well-motivated alternative to the standard Searlean account of directive insincerity. The theory is based on an empirically plausible and fruitful model of human communicative capacities—so-called pragmatic competence—namely the broadly Gricean or intentionalist program in pragmatics. I have already argued that the standard account fails to respect plausible symmetries between assertives and directives. In what follows, I address one objection and provide an additional reason to prefer the Gricean view to the Searlean.

7 Trusting the Speaker

Despite what I have argued, one might wonder why exactly it matters that the attitude produced by a directive does not match any of the speaker's own attitudes. More specifically, in placing our trust in the speaker, by deciding to comply with what we are directed to do, it seems like we rely on their intention that we do something. And insincerity, here, appears to be connected to the potential for misplaced trust; the speaker turns out to lack the attitude we were relying on.⁸

As mentioned briefly before (§5), I think the perspective of the addressee can obscure important distinctions. It is quite true that the addressee relies on the speaker in the way suggested but, still, it is important to distinguish the roles of different attitudes in grounding trust. Also, insincerity is a property of the speaker's act, regardless of how the addressee reacts. Focus on assertion first. When I assert *p* to you sincerely and truthfully, I have both the effective intention of producing an *A(p)*-state in you, and the signaling intention that you

⁸Thanks to a reviewer for this journal for raising the questions relevant to this section.

recognize my effective intention. In forming the $A(p)$ -state you correspondingly rely on two things. First, the fact that I am in the $A(p)$ -state myself (I am sincere). Second, the fact that I genuinely intend to produce the $A(p)$ -state in you (I am genuine). I have proposed to call this an authentic speech act.

So, I might express $A(p)$ ironically, without intending that you thereby form that belief. Even if I do believe p , you will not or should not rely on me to start believing p . This is because I was not genuine in expressing $A(p)$. Thus, your trust in me is of a different sort in these cases, and we should not lump these two notions together under the heading of ‘insincerity.’ That is to say, I am insincere when I don’t believe what I assert, but I am doing something different when I do not intend you to believe what I make as if to assert.

Now we can come back to directives. It is true that we rely on the speaker’s attitude of production in forming our own intention. This means that we are relying on the speaker’s effective intention, which embeds a certain other intention. So, on my view, we are relying on the speaker’s genuineness: we form our own intention in light of our confidence that the speaker genuinely wants us to form that intention. Yet, because self-talk is a real phenomenon which our theories should account for, there will be two types of reliance for directives too.

Consider the example of the meeting once more. If I realize suddenly that it is now 3pm, and that my meeting was supposed to start at 3pm, I might say to myself: *Go now!* In an ‘authentic’ case, I will rely on my own intention as well as my intention to produce that intention. To make the possibility of failure more obvious, let us introduce delay, like before. I send myself an email which is supposed to show up in my inbox at 3pm. It says: *Go to the meeting now!*

There is a sense in which I will rely on my own intention to go to the meeting. Thus I rely on the attitude produced rather than the attitude of production. If I have no intention to go to the meeting in making the utterance—writing the email—my reliance as addressee is spurious to that extent. This holds whether or not I did have the intention to produce the intention to go to the meeting. Further, the spuriousness in question is well-described as insincerity. In this case I may successfully fool my later self into thinking that my earlier self intended to go to the meeting, and this is what makes the utterance insincere. It is not made sincere, holding everything else fixed, by assuming that my earlier self was *genuine*. That is, just like in the case of assertion, speakers are not sincere merely by having the effective intention. And, certainly, introducing the delay is optional; the same distinctions will apply in principle without it.

We are considering a case where my later self trusts my earlier self to have really had the intention to go to the meeting. That is to say, my later self forms

the intention to go, partly on the basis of the belief that my earlier self had this first-order intention. If so, my later self's reliance parallels exactly the reliance of another on my asserted belief. Thus my later trust is spurious in virtue of being based on insincerity but not in virtue of being based on faking, as these terms are understood here. Admittedly, this becomes even more complicated if we include the so-called 'third clause' (sometimes known as the 'Gricean' or 'directive' intention) with the effective and signaling intentions (for discussion, see García-Carpintero 2001: 102–103; Neale 1992: 547–549; Sperber & Wilson 2015). For this clause says that the speaker intends the addressee to form the relevant attitude *partly on the basis* of recognizing the effective intention. In fact I think the third clause should be included and, so, that the complications are welcome (see, e.g., Unnsteinsson 2022b, ch. 6). For current purposes, however, we can assume that the third clause is optional. I conclude that the relevant distinctions between how different attitudes ground differences in the addressee's reliance, are available for both assertives and directives. And, to see that this is so, all we need to do is to take self-talk seriously.

8 Some Predictions

Assuming that the Gricean theorist is correct to insist that there are two pro-attitudes involved in directives—roughly, the attitude produced and the attitude of production—it can be shown that the standard account is implausible. Basically, we are owed an explanation why one of these attitudes is chosen as the one expressed rather than the other, and the choice will seem arbitrary. On the standard account, apparently, the attitude of production is chosen because it is indeed one which is available to the speaker as a first-order attitude. The speaker, in their own mind, forms the effective intention itself. But this will not do as a reason; everyone already accepts that speakers can express and communicate attitudes which they do not have, for example in paradigm insincere assertion, expressing the belief that *p* while not believing that *p*. So, we are owed an account of why the speaker expresses the second-order attitude, embedding the addressee's intention, rather than the first-order attitude itself (the produced attitude). Assertion is simpler because the attitude expressed and the attitude produced may be identical, it is assumed, so the choice seems unnecessary. But, as we have seen, if we take self-talk seriously, we must also take seriously the idea that speakers can sometimes express *their own* first-order intentions, in issuing directives, as well as any second-order desire for the addressee's compliance.

On the Gricean framework, directive competence is the capacity to utter imperative clauses with a complex communicative intention which embeds the addressee's intention to act. The corresponding interpretive competence, for directives, is the capacity to recognize speakers' communicative intentions when they properly embed prescriptions rather than propositions. The paradigm is a directive, like (1), addressed to someone else, with a literal meaning of the form $I(\pi)$. This identifies the produced attitude with what the speaker literally means. Even in such a standardized use of the directive, what the speaker literally means cannot correspond to any of their first-order mental states. This state is not available to the speaker, consciously or unconsciously. That is to say, the speaker cannot form, in their own mind, the mental state which, occurring in the mind of the addressee, constitutes the satisfaction of the meaning of the directive utterance, if the meaning is an attitude like $I(\pi)$. Thus the speaker must stand in some *other* relation to that attitude, presumably a higher-order or metarepresentational relation.

It follows that any competent user of imperatives as directives, addressing others, must distinguish between the attitude they *mean* and their own attitude in producing the utterance. In the Gricean terminology introduced here, they must distinguish between the state they intend to produce and their own mental state in performing an act whereby they attempt to produce that state. And this does not only apply to 'insincere' directives, but also paradigm ones. In the case of assertion, by comparison, some philosophers have held that sincerity is the default setting and that insincere speech—requiring some sensitivity to the distinction between what is asserted and what is believed—is best understood as pretended assertion (see, e.g., Fricker 1994, Millikan 2017, Moran 2005, Sellars 1969, Williams 2002, Rosenthal 1989). On this view effective intentions could, for all we know, be a late development in a speaker's pragmatic competence. No such move is available for directives, because even paradigm or 'sincere' utterances require the speaker's cognitive sensitivity to the very same distinction between produced attitudes and attitudes of production.

This indicates that directive competence is not acquired by being in some first-order state and then becoming able to express or give voice to that state directly, for the state is not available to the speaker as a first-order state. Speakers, even from a very young age, must then be sensitive to the distinction between the state they intend to produce and any state of their own, potentially long before they acquire the competence to issue insincere assertions. To acquire this latter competence, however, is simply to apply the same distinction to a different kind of speech act, namely assertives. So, plausibly, if there is a period in

the acquisition of pragmatic competence where speakers cannot issue insincere assertions, only sincere ones, it is because the effective intention, for some reason, embeds a state of the speaker when what is meant is a proposition rather than a prescription during that period. The reason is not that the speakers lack effective intentions, because they are already needed to explain directive competence. It is reasonable to suppose, then, that effective intentions should be generalized to cover both assertive and directive competence, for it is more likely that the same cognitive mechanism is responsible for both, rather than two distinct mechanisms. By the same token, the mechanism for effective intentions should be postulated for both paradigm and non-paradigm utterances.

As mentioned before, research in language acquisition suggests that children can acquire the imperative clause type very early, possibly earlier than other clause types (Aikhenvald 2010, 325–330). If effective intentions are necessary for the acquisition of directive competence, as I have tried to argue, it follows that children can form effective intentions before acquiring the capacity to issue insincere assertions. It would then follow, moreover, that children have the requisite capacity for metarepresentation before or during language acquisition, as predicted by Gricean models (Scott-Phillips 2015). The prediction has long been controversial, but appears to be on sound footing. Basically, directives require speakers to engage in sophisticated reversals in perspective which may appear unnecessary in the competent use of assertives, at least until we reach insincere assertives. To satisfy a directive the competent addressee must be able to activate a specific mental state, while recognizing that it is not a state of the speaker, except insofar as its activation in the addressee is what the speaker wants. To use a directive competently, however, one must be able to reverse these two perspectives. Paying attention to directives and authenticity in speech seems to suggest that effective intentions are necessary for the development of pragmatic competence. The standard account of directive insincerity is, instead, expressivist; the idea is to find a first-order attitude which the speaker can express directly. As we have seen, however, the expressivist perspective fails to do justice to insincerity in self-directed speech.

9 Conclusion

If I sincerely remind myself that I should go now to the meeting I had planned at 3pm, because 3pm is now, it seems like I must intend to attend the meeting. That is to say, my intention to attend seems to determine the extent of my sincerity

in speaking. This can be gauged, I have argued, by considering which changes would make my utterance insincere. For example, I might have no intention to attend. I might even have a non-negotiable policy of never attending meetings scheduled at 3pm. If, in that situation, I say to myself, ‘Go now!’, it seems like I am insincere to myself. My reason for making the utterance could be to make some eavesdropper believe that I’m getting ready to go, while I am not.

If this is right, what I have called the ‘produced’ attitude can sometimes determine the sincerity status of a directive utterance, namely when the attitude and the utterance are self-directed. This is a good reason, I believe, to accept the Grice-inspired distinction between attitudes of production and attitudes produced and, furthermore, to apply this distinction to any type of speech act. Moreover, by the failure of the standard account of insincerity, we are motivated to look for an alternative. I have argued that there is such an alternative view, which relativizes sincerity and related properties—authenticity, genuineness—to particular cognitive effects as well as to the distinction between self and other.

Bibliography

- Aikhenvald, A., 2010. *Imperatives and commands*. Oxford University Press.
- Anscombe, G. E. M., 1957. *Intention*. HUP.
- Bach, K., 2005. “Context *ex machina*.” Z. G. Szabó (ed.), *Semantics versus pragmatics*, Clarendon, pp. 15–44.
- Bach, K. & Harnish, R., 1979. *Linguistic communication and speech acts*. MIT Press.
- Braun, D., 2011. “Implicating Questions.” *Mind and Language*, 26(5):574–595.
- Camp, E., 2012. “Sarcasm, pretense, and the semantics/pragmatics distinction.” *Noûs*, 46(4):587–634.
- Carruthers, P., 2018. “Basic Questions.” *Mind and Language*, 33(2):130–147.
- Carson, T. L., 2006. “The Definition of Lying.” *Noûs*, 40(2):284–306.
- , 2010. *Lying and Deception: Theory and Practice*. Oxford University Press.
- Carston, R., 2002. *Thoughts and utterances*. Blackwell.
- Chan, T. & Kahane, G., 2011. “The Trouble with Being Sincere.” *Canadian Journal of Philosophy*, 41(2):215–234.
- Charlow, N., 2014. “Logic and Semantics for Imperatives.” *Journal of Philosophical Logic*, 43(4):617–664.
- , 2018. “Clause-Type, Force, and Normative Judgment in the Semantics of

- Imperatives." D. Fogal, D. W. Harris, & M. Moss (eds.), *New Work on Speech Acts*, Oxford University Press, pp. 67–98.
- Clark, H. H., 1992. *Arenas of language use*. University of Chicago Press.
- Condoravdi, C. & Lauer, S., 2012. "Imperatives: Meaning and illocutionary force." *Empirical issues in syntax and semantics*, 9:37–58.
- Currie, G., 2006. "Why Irony is Pretence." S. Nichols (ed.), *The Architecture of the Imagination: New Essays on Pretence, Possibility, and Fiction*, Clarendon Press, pp. 111–133.
- Eriksson, J., 2011. "Straight Talk: Conceptions of Sincerity in Speech." *Philosophical Studies*, 153(2):213–234.
- Fallis, D., 2012. "Lying as a Violation of Grice's First Maxim of Quality." *Dialectica*, 66(4):563–581.
- Fodor, J. & Lepore, E., 2004. "Out of context." *Proceedings and Addresses of the American Philosophical Association*. vol. 78, pp. 77–94.
- Fricker, E., 1994. "Against Gullibility." A. Chakrabarti & B. K. Matilal (eds.), *Knowing from Words*, Kluwer Academic Publishers, pp. 125–161.
- Friedman, J., 2013. "Question-Directed Attitudes." *Philosophical Perspectives*, 27(1):145–174.
- García-Carpintero, M., 2001. "Gricean rational reconstructions and the semantics/pragmatics distinction." *Synthese*, 128(1-2):93–131.
- Gregory, D., 2016. "Inner Speech, Imagined Speech, and Auditory Verbal Hallucinations." *Review of Philosophy and Psychology*, 7(3):653–673.
- Grice, P., 1968. "Utterer's Meaning, Sentence-Meaning, and Word-Meaning." *Foundations of Language*, 4(3):225–242.
- , 1969. "Utterer's meaning and intention." *The Philosophical Review*, 78(2):147–177. Repr. in Grice (1989), pp. 86–116.
- , 1975. "Logic and conversation." *Syntax and Semantics*, 3:41–58. Repr. in Grice (1989), pp. 22–40.
- , 1989. *Studies in the way of words*. HUP.
- , 1991. *The Conception of Value*. Clarendon Press.
- Hardy, J., 2006. "Speaking clearly: A critical review of the self-talk literature." *Psychology of Sport and Exercise*, 7(1):81–97.
- Hare, R. M., 1952. *The Language of Morals*. Oxford Clarendon Press.
- Harris, D. W., 2021. "Imperative Inference and Practical Rationality." *Philosophical Studies*, (4):1065–1090.
- , forthcoming a. "Intentionalism and Bald-Faced Lies." *Inquiry: An Interdisciplinary Journal of Philosophy*.
- Heal, J., 1977. "Insincerity and Commands." *Proceedings of the Aristotelian Society*,

- 77:183–201.
- Kaplan, D., 2004. “The Meaning of ‘Ouch’ and ‘Oops’.” *Howison Lecture in Philosophy*. [Http://eecoppock.info/PragmaticsSoSe2012/kaplan.pdf](http://eecoppock.info/PragmaticsSoSe2012/kaplan.pdf).
- Kaufmann, M., 2012. *Interpreting imperatives*, vol. 88. Springer Science & Business Media.
- Keiser, J., 2016. “Bald-Faced Lies: How to Make a Move in a Language Game Without Making a Move in a Conversation.” *Philosophical Studies*, 173(2):461–477.
- Langland-Hassan, P. & Vicente, A., 2018. *Inner Speech: New Voices*. Oxford University Press.
- Lepore, E. & Stone, M., 2015. *Imagination and convention: Distinguishing grammar and inference in language*. OUP.
- Ludwig, K. & Boisvert, D., 2006. “Semantics for Non-Declaratives.” E. Lepore & B. C. Smith (eds.), *The Oxford Handbook of Philosophy of Language*, Oxford University Press.
- Maitra, I., 2018. “Lying, Acting, and Asserting.” E. Michaelson & A. Stokke (eds.), *Lying: Language, Knowledge, Ethics, and Politics*, Oxford University Press, pp. 65–82.
- Michaelson, E. & Stokke, A. (eds.), 2018. *Lying: Language, Knowledge, Ethics, and Politics*. Oxford University Press, Oxford.
- Millikan, R. G., 2017. *Beyond Concepts: Unicepts, Language, and Natural Information*. Oxford University Press.
- Moran, R., 2005. “Problems of Sincerity.” *Proceedings of the Aristotelian Society*, 105(3):341–361.
- Neale, S., 1992. “Paul Grice and the philosophy of language.” *Linguistics and Philosophy*, 15(5):509–559.
- , 2016. “Silent reference.” G. Ostertag (ed.), *Meanings and other things: Themes from the work of Stephen Schiffer*, OUP, pp. 229–344.
- Portner, P., 2004. “The Semantics of Imperatives within a Theory of Clause Types.” K. Watanabe & M. Pollack (eds.), *Proceedings of SALT 14*. Cornell University Linguistics Department, CLC Publications.
- , 2007. “Imperatives and Modals.” *Natural Language Semantics*, 15(4):351–383.
- , 2018. *Mood*. Oxford University Press.
- Ridge, M., 2006. “Sincerity and Expressivism.” *Philosophical Studies*, 131(2):487–510.
- Roberts, C., 2018. “Speech Acts in Discourse Context.” D. Fogal, D. W. Harris, & M. Moss (eds.), *New Work on Speech Acts*, Oxford University Press, pp. 317–359.

- Rosenthal, D., 1989. "Postscript to "Intentionality"." *Rerepresentations: Readings in the philosophy of mental representation*, Kluwer Academic Publishers, pp. 341–344. Repr. in Rosenthal (2005), pp. 100–102.
- , 2005. *Consciousness and mind*. OUP.
- Sadock, J. M. & Zwicky, A. M., 1985. "Speech Act Distinctions in Syntax." T. Shopin (ed.), *Language Typology and Syntactic Description I: Clause Structure*, Cambridge University Press, pp. 155–196.
- Saul, J. M., 2012. *Lying, Misleading, and What is Said: An Exploration in Philosophy of Language and in Ethics*. Oxford University Press.
- Schiffer, S., 1972. *Meaning*. OUP.
- Schwartz, J., 2020. "Saying 'Thank You' and Meaning It." *Tandf: Australasian Journal of Philosophy*, 98(4):718–731.
- Scott-Phillips, T., 2015. *Speaking our minds: Why human communication is different, and how language evolved to make it special*. Palgrave Macmillan.
- Searle, J., 1969. *Speech acts: An essay in the philosophy of language*. CUP.
- , 1979. *Expression and meaning*. CUP.
- Sellars, W., 1969. "Language as Thought and as Communication." *Philosophy and Phenomenological Research*, 29(4):506–527.
- Siebel, M., 2002. "What Is an Illocutionary Point?" G. Grewendorf & G. Meggle (eds.), *Speech Acts, Mind, and Social Reality: Discussions with John R. Searle*, Springer, pp. 125–139.
- Simons, M., 2017a. "Local Pragmatics in a Gricean Framework." *Inquiry: An Interdisciplinary Journal of Philosophy*, 60(5):466–492.
- , 2017b. "Local Pragmatics in a Gricean Framework, Revisited: Response to Three Commentaries." *Inquiry: An Interdisciplinary Journal of Philosophy*, 60(5):539–568.
- Sorensen, R., 2007. "Bald-Faced Lies! Lying Without the Intent to Deceive." *Pacific Philosophical Quarterly*, 88(2):251–264.
- Sperber, D. & Wilson, D., 1986. *Relevance: Communication and cognition*. Blackwell.
- , 1995. *Relevance: Communication and cognition*. Blackwell, 2 edn.
- , 2015. "Beyond speaker's meaning." *Croatian Journal of Philosophy*, 15(44):117–149.
- Starr, W. B., 2020. "A Preference Semantics for Imperatives." *Semantics and Pragmatics*, 20.
- Stokke, A., 2014. "Insincerity." *Noûs*, 48(3):496–520.
- , 2018. *Lying and Insincerity*. Oxford, UK: Oxford University Press.
- Unnsteinsson, E., 2022a. "The Social Epistemology of Introspection." *Mind and*

- Language*, Early View.
- , 2022b. *Talking About*. Oxford University Press.
- Vranas, P. B. M., 2008. “New Foundations for Imperative Logic I: Logical Connectives, Consistency, and Quantifiers.” *Noûs*, 42(4):529–572.
- Whitcomb, D., 2010. “Curiosity Was Framed.” *Philosophy and Phenomenological Research*, 81(3):664–687.
- Williams, B., 2002. *Truth and Truthfulness: An Essay in Genealogy*. Princeton University Press.
- Wilson, D. & Sperber, D., 2012. *Meaning and relevance*. CUP.
- Winsler, A., Fernyhough, C., & Montero, I., 2009. *Private Speech, Executive Functioning, and the Development of Verbal Self-Regulation*. Cambridge University Press.
- Zanuttini, R., Pak, M., & Portner, P., 2012. “A syntactic analysis of interpretive restrictions on imperative, promissive, and exhortative subjects.” *Natural Language & Linguistic Theory*, 30(4):1231–1274.