

Is There a Meaning-Intention Problem?

JESSE RAPPAPORT

The Graduate Center, City University of New York, New York, USA

Stephen Schiffer introduced the “meaning-intention problem” as an argument against certain semantic analyses that invoke hidden indexical expressions. According to the argument, such analyses are incompatible with a Gricean view of speaker’s meaning, for they require speakers to refer to things about which they are ignorant, such as modes of presentation. Stephen Neale argues that a complementary problem arises due to the fact that speakers may also be ignorant of the very existence of such aphonetic expressions. In this paper, I attempt to articulate the assumptions that support the meaning-intention problem. I argue that these assumptions are incompatible with some basic linguistic data. For instance, a speaker could have used a sentence like “The book weighs five pounds” to mean that the book weighs five pounds on Earth, even before anyone knew that weight was a relativized property. The existence of such “extrinsic parameters” undermines the force of the meaning-intention problem. However, since the meaning-intention problem arises naturally from a Gricean view of speaker’s meaning and speaker’s reference, the failure of the argument raises problems for the Gricean. I argue that the analysis of referring-with offered by Schiffer, and defended by Neale, is defective.

Keywords: Grice, speaker meaning, reference, hidden indexical, meaning-intention problem.

1. *What Is the Meaning-Intention Problem?*

1.1. *The Meaning-Intention Problem: Background*

When confronted with some recalcitrant data, it is common practice in linguistic theorizing to invoke hidden representational structure. Hidden structures abound in syntactic theory, but they are common in semantics as well. If one is unable to generate the intuitively correct truth conditions for a sentence simply by assigning semantic values to the overt parts, one can often resolve this problem by appealing to covert variables or modifiers.

For example, it is notoriously hard to account for truth value judgments regarding belief attributions without invoking such hidden structure:

- (1) Alice believes that George Eliot is a man.
- (2) Alice believes that Mary Ann Evans is a man.

Let us suppose that Alice has heard the names ‘George Eliot’ and ‘Mary Ann Evans’ before, and she thinks they denote two distinct English authors. Then it is plausible to suppose that (1) and (2) might receive different truth values. But a naïve, extensional analysis of proper names would make it difficult to account for this difference. This is in part the motivation for Fregean theories of proper names.

On a Fregean analysis, expressions appearing in intensional contexts do not denote their reference, but instead denote their (customary) *sense*. Since ‘George Eliot’ and ‘Mary Ann Evans’ differ in sense, their occurrences in (1) and (2) differ in reference, and hence (1) and (2) may differ in truth value. Thus, although a classical Fregean analysis does not require positing any hidden *structure* in a belief report, it does imply that the proposition meant contains *senses*, entities of which speakers may lack a concept.

A competing account can be found in the writings of neo-Russellian theorists such as Nathan Salmon (1986). These theorists differ from the Fregeans in holding that the semantic value of a proper name is exhausted by its referent, however they explain the apparent failure of substitutivity by positing substantial hidden complexity in belief attributions. In particular, whereas a naïve view might hold that *believes* expresses a two-place relation between a believer and a proposition believed, Salmon proposes that belief attributions in fact express a three-place relation between a believer, a proposition believed, and a propositional *guise* under which the believer accepts the proposition:

- (3) *For some guise x*, Alice grasps that George Eliot is a man by means of *x* and BEL(Alice, that George Eliot is a man, *x*).

According to this view, the semantic analysis of both (1) and (2) is (3). Thus, both sentences are literally true. However, (2) is highly misleading, and pragmatically conveys something false (e.g. that Alice would assent to ‘Mary Ann Evans is a man’). Thus, we are simply judging the sentence according to what is pragmatically conveyed, rather than its semantic content. This view also allows us to account for sentences like,

- (4) Alice believes that Mary Ann Evans is not a man,

as follows:

- (5) *For some x*, Alice grasps that George Eliot is not a man by means of *x* and BEL(Alice, that George Eliot is not a man, *x*).

Alice may accept (BEL) that George Eliot/Mary Ann Evans is a man under some guises, and reject it (or accept its negation) under other guises. This elegant analysis allows one to preserve a simplified, Mil-

lian view of proper names, while accommodating the cases of substitution failure in a fairly systematic way.

However, the neo-Russellian analysis of belief attribution also requires the theorist to posit entities—guises—of which the speaker is ignorant, in the semantic analysis of ordinary belief reports. In addition, it implies that the logical structure of a belief attribution is more complex than an ordinary speaker might recognize. These facts are the basis of a persuasive argument against such theories, first developed by Stephen Schiffer (1992). This argument, known as the *meaning-intention problem*, claims that such analyses cannot be correct, because speakers do not, and cannot, mean what such analyses require them to mean.

Schiffer develops his argument in response to the *hidden-indexical theory* of belief attributions. This view, similar to Salmon's, analyzes belief reports as follows:

- (6) Ralph believes that Fido is a dog.
- (7) (*For some m*) (\square^*m & B(Ralph, <Fido, doghood>, *m*))

On this view, (6) is analyzed as the statement that Ralph believes that Fido is a dog under some mode of presentation (MOP) *m*, where *m* satisfies some contextually specified constraint on MOPs \square^* .

“The meaning-intention problem” in fact covers a number of related concerns raised by Schiffer for this type of analysis:

- (a) The Awareness Problem—If the hidden-indexical theory of belief reports is true, then speakers lack full, conscious awareness of what they mean, for most speakers have no idea that they are referring to a MOP-property: “Thus, if the hidden-indexical theory is correct, then [the speaker] has no conscious awareness of what she means, or of what she is saying... and this is a *prima facie* reason to deny that she means what the theory is committed to saying she means” (Schiffer 1992: 514).
- (b) The Cognitive-Resources Problem—If the hidden-indexical theory is true, then what a speaker means may involve entities about which the speaker is totally ignorant (such as MOP-properties).
- (c) The Specificity Problem—If the hidden indexical theory is true, then there must be some *particular* property \square^* of MOPs that the speaker is referring to. But it's doubtful that the speaker's intentions serve to pick out any such particular property (the speaker lacks “specifying intentions,” we may say).

On the assumption that speakers *do*, generally, have full, conscious awareness of their own speech and thought contents, and that speaker-meaning requires a speaker to have a conceptual grasp of the particular entities and properties that comprise the contents of their speech acts, the meaning-intention problem raises serious doubts about the hidden-indexical theory of belief reports. Similar doubts would arise for any theory that posits, in the semantic analysis of some sentence, entities,

properties, or structure about which the speaker (a) lacks conscious awareness, (b) is totally ignorant, or (c) lacks specifying intentions.

Schiffer (1996) deploys this argument against another prominent contextualist theory that invokes hidden indexicals: contextualist theories of knowledge claims. Such theories are in part an attempt to respond to the fact that a sentence like ‘I know that I have hands’ might be judged true in an ordinary context, but false in the context of a discussion about Descartes’s evil demon. Contextualists claim that knowledge claims are tacitly relativized to a *standard of knowledge*, which affects the kind of evidence needed to support a knowledge claim in a given context. Thus, ‘I know that I have hands *relative to a low standard of knowledge*’ might be true, while ‘I know that I have hands *relative to a high standard of knowledge*’ might be false.

Schiffer sees two problems with this approach. First, it again implies that speakers are not fully aware of what they mean, think, or say: “But no ordinary person who utters ‘I know that *p*’, however articulate, would dream of telling you that what he meant and was implicitly stating was that he knew that *p* relative to such-and-such standard” (Schiffer 1996: 326–327). Second, it implies that speakers’ conception of their own speech contents may not only be *incomplete*, but may be seriously mistaken. For the contextualist hopes to explain why the following kind of argument *appears* to be sound, even though the conclusion seems incorrect:

- (A) I don’t know that I’m not a brain-in-a-vat.
- (B) I know that I have hands only if I know that I’m not a brain in a vat.
- (C) Therefore, I don’t know that I have hands.

The response is that (C) strikes us as false because we are prone to interpret it as:

- (C*) Therefore, I don’t know that I have hands *relative to a low standard of knowledge*,

even when the context dictates that we ought to interpret it as:

- (C**) Therefore, I don’t know that I have hands *relative to a high standard of knowledge*.

The problem, according to Schiffer, is that it’s highly implausible to suppose that speakers could be so confused about the content of their own speech acts: “It’s as though a fluent, sane, and alert speaker, who knows where she is, were actually to assert the proposition that it’s raining in London, when she mistakenly thinks she’s asserting the proposition that it’s raining in Oxford” (Schiffer 1996: 326).

1.2. *Assumptions Behind the Meaning-Intention Problem*

As we have seen, the meaning-intention problem attempts to cast doubt on contextualist semantic theories by arguing that we can’t take

speakers to mean what such theories require them to mean. However, the argument is deployed with different emphasis in different places. Thus, it is worthwhile to clearly articulate the assumptions upon which such an argument rests:

(Awareness) If a speaker, *S*, performs an utterance, and means thereby that *p*, and *x* is a constituent of the proposition *p*, then *S* must be consciously aware that *x* is part of what she means.

(Cognitive Resources) If a speaker, *S*, performs an utterance, and means thereby that *p*, and *x* is a constituent of the proposition *p*, then *S* must have a concept of *x*, whether consciously accessible or not.

(Specificity) If a speaker, *S*, performs an utterance, and means thereby that *p*, and *x* is a constituent of the proposition *p*, then *S* must have specifying intentions that uniquely determine reference to *x*, rather than any other nearby candidates that might appear to serve just as well in the context.

Here, I adopt a structured-proposition view of propositional content for ease of exposition. Saying that *x* is a constituent of the proposition *p* is simply an attempt to capture the intuitive idea that *p* is, in some sense, “about” *x*. If a speaker utters ‘It’s raining’ and means thereby that it’s raining in London, then she has said something about London. The truth conditions of her utterance are sensitive to how things are in London.

Although these principles are distinct, they all attempt to use the fact that meaning is an intentional act performed by speakers as a way to constrain what counts as an acceptable semantic analysis. In Section II, I argue that these principles are not sound, and that the meaning-intention problem, as stated, is far too powerful.

1.3. *The Aphonic-Intention Problem*

Stephen Neale (2016) defends Schiffer’s meaning-intention problem, and claims there is an additional problem related to the use of aphonic referring expressions. The problem is based on the Gricean analysis of what is it to refer *with* an expression:

(RW) In uttering *x*, *S* referred to *o* *with* (or *using*) *e*, relative to its *i*-th occurrence in *x*, iff for some audience *A* and relation *R*, *S* intended *A* to recognize that $R(e, x, i, o)$ and, at least partly on the basis of this, that *S* referred to *o* in uttering *x*.

R, in this definition, is, effectively, the *inference-base* feature (cf. Schiffer 2017) of the expression *e*, that is, the property that *S* believes *e* has (relative to its position *i* in the sentence *x*), such that the audience will recognize that *S* is referring to *o* partly on the basis of recognizing this feature. Often, this feature is simply the fact that *e* is conventionally used to refer to *o*.

This definition of *referring-with* invokes the more basic notion of *speaker-referring*, which Neale follows Schiffer in defining as:

(SR) In f-ing, *S* referred to *o* iff what *S* meant by f-ing is an *o*-dependent proposition (a singular proposition that has *o* as a constituent).

This definition, in turn, analyzes speaker-referring in terms of the more basic notion of *speaker's meaning*, and hence, together these definitions allow us to explain what it is for a speaker to refer *with* an expression in basic Gricean terms.

What's important, for the present discussion, about this Gricean analysis of referring-with is that it entails that if a speaker uses *e* to refer to *o*, then the speaker must have an intention that *has the expression e as part of its content*. It's not simply that, in referring to *o*, a speaker must employ the expression *e* in actualizing her intention, the way she must perform an alveolar stop in pronouncing the word *to*, something which most English speakers are capable of, whether or not they have a concept of *alveolar stop*. For in general, it isn't the case that if actualizing some intention requires a subject to utilize some capacity, then she must be aware of possessing that capacity. However, due to the nature of communicative intentions, if a speaker uses some expression *e* to refer to *o*, then she must intend for the audience to recognize *that e has some inference-base feature*, and thus, she must be capable of forming intentions whose content involves the expression *e*. This implies that referring with some expression *e* requires the subject to have some level of conceptual grasp of the expression *e* itself.

The problem that Neale raises is that this consequence is extremely dubious for cases of aphonic referring expressions, since most speakers have no idea that such expressions exist. Thus, there is a meaning-intention problem not only with respect to the entities that a speaker refers *to*, but also with respect to the entities that speakers refer *with*: "So an implicit reference theory according to which speakers refer aphonically to mode of presentation types faces a compound problem: the theory has ordinary speakers referring *to things* they don't know about *with things* they don't know about" (Neale 2016: 154).

This variation of the meaning-intention problem appears to rely on the following assumption, which is a fairly straightforward consequence of (RW):

(Syntactic Knowledge) If a speaker *S* refers to some entity *o* with some expression *e*, then *S* must be able to form an intention that has *e* as part of its content, and hence must "know about" *e*, in some sense.

I will argue that this assumption is incompatible with the data.

2. *Extrinsic Parameters*

There are a number of examples that appear to contradict the proposed principles supporting the meaning-intention problem. They seem to show quite clearly that a speaker *S* can express a proposition *p* that is about, or concerns, some entity *e*, even though *S* either lacks a concept of *e*, *contra* (Cognitive Resources), or simply lacks conscious awareness of expressing a proposition about *e*, *contra* (Awareness).

- (8) This book weighs five pounds.
- (9) The flash and the bang happened at the same time.
- (10) It's summer.

Each of these sentences expresses a proposition whose truth value depends on the state of some entity that is not explicitly mentioned in the sentence itself:

- (8*) This book weighs five pounds (*on Earth*).
- (9*) The flash and the bang happened at the same time (*relative to Earth as a frame of reference*).
- (10*) It's summer (*in the Northern hemisphere*).

Most educated speakers know that weight is something that can actually vary depending on what large body of mass you happen to be standing on. And many (though perhaps not most) speakers know that simultaneity is not an absolute relation, but rather that two events are simultaneous only with respect to a reference frame.

What's important for our discussion is that despite the fact that many educated speakers are aware of such relativization, *many competent speakers are not*. Indeed, for large periods of human history, *no competent speaker* grasped that weight is relativized or had any idea what a reference frame was. This did not in any way prevent them from communicating complete thoughts in uttering sentences like (8)–(10). Therefore, *contra* (Cognitive Resources), it is not the case that if a speaker expresses some proposition *p* that contains some entity *o*, then she must have some conceptual grasp of *o*. *A fortiori*, she need not have conscious awareness of such a concept either, *contra* (Awareness).

Considering sentence (10), even if we assume that all, or almost all, competent speakers know that what season it is depends on what hemisphere you are in, it is entirely plausible that a speaker might utter (10) and not be thinking about hemispheres at all, i.e., not have any awareness of saying something that is about a hemisphere. This again contradicts (Awareness). Even for those cases in which speakers have the requisite concepts, they may not be aware of invoking those concepts in performing their utterance.

Neale (2016: 160ff.) raises similar examples, referring to such cases as involving “extrinsic parameters,” but fails to draw any substantial conclusions from them. Instead, he focuses on the disanalogies between these cases and MOPs, in order to show that whatever solace the hidden indexical theorist hopes to find in such cases does not help the hidden indexical theory of belief reports. Examples (8)–(10) relate to

factors *external* to us... about which we may be ignorant but about which we may acquire knowledge and thereby easily refine our linguistic behaviour. Mode of presentation types are not like this at all. They are supposed to be things *under* which beliefs are *had*, and learning about their existence and a great deal of information about their roles in theories of language and mind doesn't even put theorists in a position to articulate the truth conditions of the propositions they actually express on given occasions us-

ing belief sentences if the hidden-indexical theory of belief reports is true. (Neale 2016: 163)

Thus, unlike with MOPs, “once speakers *learn* about time-zones, hemispheres and rest-frames, and learn a few additional words, they can easily describe the parameters relevant to the truth or falsity of what they are saying...”

First of all, this seems doubtful. For instance, I know that whether or not two events are simultaneous depends on a reference frame. However, not having much grasp of relativity theory, my knowledge of reference frames ends there. It is fair to say that I know far more about MOPs than about reference frames.

The main issue, though, is not whether MOPs are easier to grasp than reference frames, or *vice versa*. The problem is rather that examples like (8)–(10) seriously undermine the general assumptions that support the meaning-intention problem. It isn’t clear what comfort the *proponent* of the meaning-intention problem is supposed to find in the fact that even though most speakers lack a proper concept of a reference frame, and many would never “dream of telling you that what he meant and was implicitly stating” was something to do with reference frames, there are *others* who do grasp the concept, and perhaps with sufficient training the rest of the population could do so as well.

The fact is that relatively few competent speakers grasp the concept of a reference frame. Some might be able to recall that simultaneity is a relative notion, if pressed, but even this is reserved for an educated segment of the population. Nevertheless, ordinary speakers are perfectly competent with phrases like “at the same time.” This, again, shows that the principles adduced to support the meaning-intention problem cannot be sustained. There may be important differences between MOPs and reference frames, but these differences cannot be used to salvage the meaning-intention problem in its current form.

What about Neale’s syntactic analogue of the original meaning-intention problem? Is there an aphonic-intention problem?

I claim, once again by *reductio ad absurdum*, that (Syntactic Knowledge) simply cannot be supported in light of the data. Consider the following sentence of Spanish:

- (11) Quiere comer. [*He wants to eat.*]

According to standard assumptions of generative syntax (cf. Haegeman 1994: 68–69ff.), in order to comply with the Extended Projection Principle, it is argued that (11) must contain a phonologically null subject, which is typically expressed as *pro*:

- (12) *pro* quiere comer.

Pro (distinct from *PRO*) is a phonologically null pronoun, which is the subject of the main clause, and whose semantic value is determined by the speaker’s intentions. Importantly, although its existence is not supported by every syntactician, *pro* is an established posit in syntax

with independent syntactic support—in other words, it is not something that is just posited by philosophers in order to generate their desired truth conditions.

Precisely the same concerns that Neale raises for aponic indexicals in ‘Silent Reference’ would have to apply to *pro* (which is, effectively, an aponic indexical). Most speakers lack any conceptual grasp of the expression *pro*, and the speaker of (11) might positively deny that she used a tacit referring expression at all. Therefore, if Neale’s arguments are sound, we find that we must reject some basic posits of mainstream syntactic theory as well.

Regardless of whether *pro* exists or not, is doubtful that Neale would be sanguine about this consequence. He often takes philosophers to task for their tendency to make non-trivial claims about syntax purely on the basis of philosophical considerations. But that is precisely what we would have to do if we accept (Syntactic Knowledge).

On reflection, we can see that (Syntactic Knowledge) is a rather demanding principle. It effectively implies that linguistic ability and metalinguistic knowledge must proceed “in tandem”—I cannot *use* an expression to refer unless I have conceptual grasp of that expression. Even for overt expressions, however, it is not obvious that we should accept such a principle.

The problem for the Gricean is that (Syntactic Knowledge) is a fairly direct consequence of (RW). But it appears that (Syntactic Knowledge) must be rejected. Therefore, it seems that (RW) must be rejected as well.

3. *Tacit States*

One direct consequence of the preceding discussion is that (Awareness) should be abandoned. This is all for the good, for the picture of meaning that it presupposes is a strongly Cartesian one. Why should we assume that speakers *do* have privileged access to every aspect of the contents of their speech acts? This is certainly not required in order for them to have meaning-intentions, assuming that such intentions may fail to be fully conscious. And why shouldn’t this be the case? Certainly, some strong arguments would be needed to establish that intentions must be conscious; or, at least, that meaning-intentions are special in that they must be conscious. But the Gricean ought to be very cautious about the latter claim, for her entire program is based around positing a certain kind of complex intention as the basis for communication, where this intention does not simply reveal itself through introspection.

The question, then, is whether (RW) can be salvaged by interpreting the intentional verbs *intend* and *recognize* in terms of tacit states—tacit intentions and tacit recognition.

First, it is worth noting that (RW), as it stands, is inadequate—it fails to provide either necessary or sufficient conditions for referring-with:

(i) Consider the following sentence:

(13) I saw Alice_i and then she_i disappeared.

Here I have referred to Alice, and there are two expressions *with which* I referred to Alice. The problem is only the first instance can satisfy (RW). Recall (RW):

(RW) In uttering x , S referred to o *with* (or *using*) e , relative to its i -th occurrence in x , iff for some audience A and relation R , S intended A to recognize that $R(e, x, i, o)$ and, at least partly on the basis of this, that S referred to o in uttering x .

(RW) requires that if a speaker S uses e to refer to o , then S must intend her audience A to recognize that S referred to o in uttering the sentence x , and to do so at least partly on the basis of her utterance of e . But the problem is that A will recognize that S referred to Alice simply on the basis of the utterance of 'Alice.' Since "referring to Alice" really means "expressing an Alice-dependent proposition," then A will know that S referred to Alice as soon as 'Alice' is uttered (and S will know this). Therefore, A will recognize (and S will know that A will recognize) that S referred to Alice purely on the basis of the first reference to Alice—and if this provides a sufficient reason for A to believe that S referred to Alice (as indeed it does), then S cannot intend for A to recognize this even *partly* on the basis of the second reference to Alice ('she'). But, intuitively, S referred to Alice with 'she.' Therefore, (RW) does not provide a necessary condition for referring-with.

To resolve this problem, (RW) must therefore be modified along the following lines:

(RW*) In uttering x , S referred to o *with* (or *using*) e , relative to its i -th occurrence in x , iff for some audience A and relation R , S intended A to recognize that $R(e, x, i, o)$ and to recognize that $R(e, x, i, o)$ *provides a reason* to believe that S referred to o in uttering x .

This modification avoids the preceding worry because it does not require that S intend that A 's belief that S referred to o be derived *on the basis of* A 's recognition that $R(e, x, i, o)$, but simply that such recognition provide the hearer with *a basis* for arriving at such a belief (whether or not it is the basis that is in fact used). However, (RW*) does *not* require S to even be referring to o (or even to intend for A to recognize that she is referring to o)—it only makes the weaker requirement that S intend for A to recognize that S is doing something that *provides a reason* to believe that S referred to o . This seems too weak. Thus, perhaps the following will suffice:

(RW**) In uttering x , S referred to o *with* (or *using*) e , relative to its i -th occurrence in x , iff for some audience A and relation R , S intended A to recognize that $R(e, x, i, o)$ and to recognize that $R(e, x, i, o)$ provides a reason to believe that S referred to o in uttering x , *and to recognize that S referred to o in uttering x .*

(ii) The second problem with (RW) is that it does not only apply to referring expressions. Consider, e.g., a classroom in which there are two teachers, Alice and Bob, and a number of young students, one of whom is named Bob. Alice is talking to Student Bob's parents and says:

(14) Bob will be teaching the class today,

meaning thereby that Teacher Bob would be teaching the class. In uttering (14), Alice relies on the conventional meaning of *Bob* as a name for Bob, and thus, following (RW), refers to Bob with *Bob* insofar as she expects her audience to recognize that she is referring to Bob by uttering *Bob*. However, she can't expect that uttering *Bob* is *sufficient* for her audience to recognize that she is referring to Teacher Bob, and not Student Bob. Rather, it is the context of the sentence as a whole—or, in particular, the verb 'teaching'—that makes it clear which Bob Alice is referring to. Thus, the right-hand side of (RW) appears to be satisfied by the utterance of 'teaching' as well. (RW) says that a speaker *S* refers to *o* with some expression *e* just in case for some audience *A* and relation *R*, *S* intended *A* to recognize that $R(e, x, i, o)$ and, at least partly on the basis of this, that *S* referred to *o* in uttering *x*. Instantiating the variables: Alice intended her audience to recognize that 'teaching' bears some relation to Teacher Bob and, at least partly on the basis of this, that *S* referred to Teacher Bob in uttering *x*. Thus, (RW) implies that Alice referred to Teacher Bob *with* the word 'teaching.' Since, intuitively, Alice did *not* refer to Bob *with* the word 'teaching,' this implies that (RW) does not provide a sufficient condition for referring-with either.

Unfortunately, no obvious solution to this problem presents itself. The simplicity of (RW) lies in the fact that it analyzes referring-with in terms of offering reasons to believe that one is expressing an *o*-dependent proposition. But, *prima facie*, there is no reason why this condition should be satisfiable by referring expressions only, since other information in the sentence might be intended to help convey what the speaker is referring to, as well.

These problems of definition notwithstanding, the question remains whether (RW) (or (RW**)) is acceptable if one reads the intentional verbs in terms of tacit states.

Two issues must be separated. In discussing tacit states, Neale suggests that 'tacit' amounts to 'unconscious,' glossing "The Tacit States Reply" as assuming that "S is not 'consciously aware' that she means a proposition of the form..." (Neale 2016: 163). This appears to be the sense of 'tacit' used by Brian Loar (1976) as well. However, Neale also describes the view as assuming "tacit knowledge" in Chomsky's conception of the term. Chomskyan tacit knowledge is not merely unconscious, but is also functionally isolated—grammatical knowledge is not integrated into the web of belief, i.e., it is not accessible to central reasoning processes. One might assume that these two categories bear some logical relation to each other—for instance, that anything that is in the web of belief is accessible to consciousness—but this is by no means obvious.

Which of these senses is more appropriate for (RW)? (RW) must grant that speakers tacitly know, e.g., that the Spanish sentence *Quiere comer* has an unpronounced nominal expression in subject position. But it is doubtful that all competent Spanish speakers have a concept (even an unconscious one) of the expression *pro*. Rather, this “knowledge” is more akin to the kind of syntactic knowledge that comprises the language faculty—not fully conceptual, and not integrated into central reasoning. Therefore, the defender of (RW) must grant that the tacit knowledge mentioned in the definition is tacit in the *strong*, Chomskyan sense.

However, introducing tacit intentions in this manner presents a radical break from traditional Gricean thinking. For one of the apparent advantages of Grice’s theory is that it accounts for the ways in which linguistic production interacts with global reasoning abilities. Speakers intentionally produce effects in the audience that are responsive to background knowledge, features of the context, and other factors that interact with central reasoning. By claiming that referential intentions can be tacit in the Chomskyan sense (functionally isolated), the defender of (RW-T) weakens the connection between speaker’s meaning and rationality. If one pursues this approach to the end, one may find that the result is no longer distinctly Gricean.

4. *Specificity*

I have argued that examples like,

(9) The flash and the bang happened at the same time,

show convincingly that the principles underlying the various aspects of “the meaning-intention problem,” in particular (Awareness) and (Cognitive Resources) are not tenable.

If we reject (Cognitive Resources), then we must reject (Specificity) as well. The problem raised by (Specificity) is that it seems like there ought to be something about the speaker’s state of mind, or communicative intentions, that determines *which* of a number of competing alternatives is the actual semantic value of a hidden indexical. But if speakers can refer to *o* (express an *o*-dependent proposition) without even having a concept of *o*, as I claim is demonstrated by sentences like (9), then it cannot be the speaker’s specifying intentions that resolve the (Specificity) problem.

However, it is worth noting that (Specificity) is not essentially a problem about language or communication, but rather is a puzzle about thought contents. For example, consider:

(15) The book is covered with paint.

Suppose a speaker utters (15) to communicate that the addressee’s favorite book, *War and Peace*, which was left sitting on the dining room table, is covered with paint. The problem relating to (Specificity) is that *the book* does not serve to uniquely pick out the aforementioned book,

and there does not seem to be any unique candidate for completing the nominal that is the one that the speaker intended to communicate. This is what is known as an *incomplete description*.

This presents a puzzle about the speaker's *thoughts* about the book, as well. In particular, what is the conceptual content of the speaker's thought, *the book is covered with paint*? One might assume that when a speaker tokens such a thought, there is some description or other (over and above *the book*) under which she thinks of the book. But why would she token one description or the other on any given occasion? Is doing so *necessary* in order for her to entertain thoughts about the book? But surely she needn't token a *uniquely* identifying description of the book, for coming up with such a description is no simple task. But then it begins to look as though the descriptive information is not doing any cognitive work, and the thought is already complete as it is. But if that's the case, then why should an utterance like (15) not be complete as it is?

In sum, I am arguing that (a) examples like (9) show that (Specificity) must be rejected if (Cognitive Resources) is rejected, as I claim it must be; and, (b) the issue of incomplete definite descriptions is primarily an issue regarding mental contents, and only derivatively an issue pertaining to communication. Furthermore, brief consideration of the matter suggests that incomplete definite descriptions may very well exist in thought as well, which implies that the assumptions behind (Specificity) are misguided.

5. *Constraints on Semantic Theories*

If these arguments are correct, then the meaning-intention problem is based on implausible assumptions, and cannot be used as a way to refute the hidden indexical theory of belief attribution, or other hidden indexical theories. Speakers can express a proposition that is *o*-dependent even if they entirely lack the concept of *o*. However, there are significant differences between the Extrinsic Parameters Reply as applied to reference frames and as applied to MOPs and other more speculative philosophical entities.

According to Hofweber (1999), the relevant difference is that when it comes to reference frames (with respect to simultaneity) or bodies of mass (with respect to weight), the extrinsic parameter, or unarticulated constituent, is constant for all members of the speech community. Thus while people failed to realize that weight is a relativized notion, there could never be any confusion or disagreement amongst the members of the community regarding which planet was relevant to their weight statements. This allowed discourse about weight to carry on unproblematically, without the body of mass being specified by the speakers' intentions. This is because the relativization did not affect "sameness and difference (or incompatibility) of contents" amongst speakers.

The same facts do not apply to MOPs. MOPs are *not* shared amongst

all members of the community, and there *is* the tendency for confusion when making belief reports.

Hofweber concludes that the problem with the hidden indexical theory of belief reports is that it implies “not only... that speakers have no access (in the strong sense spelled out above) to the content of their utterances, but also no access to sameness, difference and incompatibility of the contents of their utterances” (Hofweber 1999: 102). However, he leaves it open why this fact is problematic.

I suspect that the problem is ill-formulated as once again pertaining to *access*. The reason why speakers must “share the referent” when it comes to genuinely unarticulated constituents (extrinsic parameters) is that testimony would break down otherwise. For two speakers might use the same sentence to express thoughts that concern distinct entities without their being aware of any difference. Thus, if one speaker asserts that sentence to another, the addressee might form a belief that has entirely different truth conditions. Over time, this discrepancy might cause problems and lead to the extinction of the use of the term.

Alternatively, through the course of such discrepancies, speakers may learn to identify in what ways the truth conditions of each speaker’s use diverge from each other, and thus may gain a conceptual understanding of the way that the sentence is relativized. This might be an important step in coming to grasp the true nature of the relation being discussed.

In sum, defenders of hidden indexical theories for belief attributions and knowledge claims cannot appeal to extrinsic parameters to supply the referents in such utterances, since the alleged referents are not constant amongst members of the community, or across contexts. The only alternative is to argue that speakers really do grasp concepts such as MOPs and their communicative intentions do serve to pick out properties of such MOPs in the required way. This would be to abandon the Extrinsic Parameters approach, but this leads back to the same problems initially raised by Schiffer and Neale.

6. *Conclusion*

The meaning-intention problem has been used by theorists in the Gricean tradition to reject certain semantic hypotheses involving hidden indexical expressions. Hidden indexical theories imply that speakers lack full awareness of their own speech and thought contents, and in some cases lack the relevant concepts entirely. This appears incompatible with the Gricean approach to language, according to which speech act content depends on the speaker’s communicative intentions. I have attempted to reverse this dialectic by showing that the assumptions that support the meaning-intention problem are incompatible with some basic data. Since these assumptions are natural ones for the Gricean to make, I have attempted to explore how this affects the project of intention-based semantics. I have argued that the definition

of referring-with that is offered by Schiffer and Neale needs to be re-considered. It fails to provide either necessary or sufficient conditions for referring with an expression, and it is incompatible with facts about *pro-drop* languages, such as Spanish.

The meaning-intention problem, as traditionally conceived, thus appears to have little force. Nonetheless, I agree with Hofweber that appeals to extrinsic parameters must be tightly constrained. The clear cases of extrinsic parameters, or unarticulated constituents, that I have appealed to (such as weight and simultaneity) involve parameters that are constant amongst the members of the speech community. MOPs and standards of knowledge do not share this feature, and thus the same considerations do not apply.

References

- Haegeman, L. 1994. *Introduction to Government and Binding: Second Edition*. Malden: Blackwell.
- Hofweber, T. 1999. "Contextualism and the Meaning-Intention Problem." In Korta, K. et al. (eds.). *Cognition, Agency and Rationality: Proceedings of the Fifth International Colloquium on Cognitive Science*. New York: Springer: 93–104.
- Loar, B. 1976. "The Semantics of Singular Terms." *Philosophical Studies* 30 (6): 353–377.
- Neale, S. 2016. "Silent Reference." In Ostertag, G. (ed.). *Meanings and Other Things: Essays in Honor of Stephen Schiffer*. Oxford: Oxford University Press: 229–342.
- Salmon, N. 1986, 1991. *Frege's Puzzle*. Atascadero, California: Ridgeview Publishing Company.
- Schiffer, S. 1992. "Belief Ascription." *The Journal of Philosophy* 89 (10): 499–521.
- Schiffer, S. 1996. "Contextualist Solutions to Scepticism." *Proceedings of the Aristotelian Society* 96 (1): 317–333.
- Schiffer, S. 2017. "Intention and Convention in the Theory of Meaning." In Hale, B., Wright, C., & Miller, A. (eds.). *A Companion to the Philosophy of Language, Second Edition*. Malden: Wiley Blackwell: 49-72.

