

Toward a Sharp Semantics/Pragmatics Distinction

Megan Henricks Stotts

This is a post-peer-review, pre-copyedit version of an article published in *Synthese*. The final authenticated version is available online at: <https://doi.org/10.1007/s11229-018-1733-8>.

Abstract:

The semantics/pragmatics distinction was once considered central to the philosophy of language, but recently the distinction's viability and importance have been challenged. In opposition to the growing movement away from the distinction, I argue that we really do need it, and that we can draw the distinction sharply if we draw it in terms of the distinction between non-mental and mental phenomena. On my view, semantic facts arise from context-independent meaning, compositional rules, and non-mental elements of context, whereas pragmatic facts are a matter of speakers' mental states and hearers' inferences about them. I argue for this treatment of the distinction by comparing it to some other extant treatments (in terms of "what is said," and in terms of the involvement of context) and then defending it against several challenges. Two of the challenges relate to possible intrusion of mental phenomena into semantics, and the third has to do with possible over-restriction of the domain of pragmatics.

1. Introduction: An Embattled Distinction

The semantics/pragmatics distinction was once considered central to the philosophy of language. Loosely, the distinction is between what words themselves mean (on the semantic side) and what particular utterances actually communicate (on the pragmatic side). Consider a letter of recommendation for graduate school in philosophy that consists of only the following: 'Anne speaks fluent English and is always punctual' (Grice 1989, p. 33). Here it seems natural to say that the semantic content of the utterance is a straightforward attribution of two positive qualities to Anne, but pragmatically, the utterance conveys something else in addition: that Anne is not qualified for graduate school.

Recently, there have been some strongly worded challenges to the semantics/pragmatics distinction's viability and importance, such as the following:

The late 90's to mid 00's played host to all manner of fashion faux pas: trucker hats, denim suits, the mullet, rap metal. Pop culture has by now moved on from most of these. Philosophy of language, on the other hand, is still largely mired in its turn-of-the millennium obsession with policing the boundary between semantics and pragmatics (Armstrong and Michaelson 2016, p. 139).

My own experience was that, over time, I simply lost my confidence in my ability to see that semantics-pragmatics line, or even to know quite what it was I was looking for. The result, as is often the way with childhood faiths, was not so much a renunciation but simply a falling

away, in the sense of ceasing to be able to see why the question mattered so much (Dever 2013, p. 105).

[T]here's no such thing as the semantics–pragmatics distinction and looking for it is a waste of time. No such distinction will do any important explanatory work.... The best solution would be for all of us to decide never to use these dreaded words ever again (Cappelen 2007, pp. 3, 20).

Here, the semantics/pragmatics distinction is portrayed as an outdated fashion trend, a childish article of faith, and a waste of time.

I find this dismissiveness understandable. Different theorists use the terms 'semantic' and 'pragmatic' in different, often irreconcilable ways, which can sometimes be a symptom of terminology that ought to be retired.¹ But in opposition to the growing movement away from the distinction, I will argue that we really do need the semantics/pragmatics distinction, and that we can draw the distinction sharply if we draw it in terms of another distinction, between non-mental and mental phenomena. On my view, semantic phenomena are a matter of context-independent expression meaning,² compositional rules, and non-mental elements of the context of utterance, whereas pragmatic phenomena are a matter of speakers' mental states and hearers' inferences about them.

In Sections 2 and 3, we'll develop criteria for a satisfactory treatment of the semantics/pragmatics distinction by discussing problems with several influential strategies for drawing it: one strategy based on Paul Grice's notion of what is said, and two based on the role of context. Along the way, we'll also see why the semantics/pragmatics distinction matters. Then, in Section 4, we'll see how drawing the semantics/pragmatics distinction in terms of the distinction between non-mental and mental phenomena satisfies our criteria. We'll discuss several challenges to Section 4's way of drawing the semantics/pragmatics distinction in Sections 5, 6, and 7.

¹ Maite Ezcurdia and Robert J. Stainton provide a helpful overview of some of the very different ways in which the terms 'semantic' and 'pragmatic' get used (2013, pp. xiii–xix).

² Context-independent expression meaning is the public meaning of an expression in a language. Below I will sometimes refer to this kind of meaning as just 'expression meaning' or 'context-independent meaning,' as a matter of emphasis.

2. The “What Is Said” Strategy

Contemporary discussion of the semantics/pragmatics distinction has its roots in Grice’s work, so we’ll begin with the strategy for drawing the distinction that is most closely linked to his writings: the “*What Is Said*” Strategy. We’ll discuss a problem the strategy faces, which will allow us to develop our first criterion for a satisfactory treatment of the distinction. Discussing Grice will also reveal why the semantics/pragmatics distinction matters.

In Grice’s terminology, in the letter of recommendation example, the claim that Anne is fluent in English and always punctual is *what is said*, whereas the claim that Anne is not qualified for graduate school is *what is implicated* (or in other words, the content of a *conversational implicature*). Once the reader grasps what was said, he is able to move from there to what was implicated by means of the assumption that the professor who wrote the letter is cooperating. A one-sentence letter of recommendation that doesn’t address the most philosophically relevant qualities of the subject of the letter initially appears highly uncooperative. But the reader will quickly realize that the professor *can* be thought of as cooperative if he (the reader) assumes that the professor believes that these are Anne’s *only* good qualities that bear any relevance to her qualifications for graduate school, and that Anne is therefore not qualified (Grice 1989, pp. 30–31, 33). The “What Is Said” Strategy construes the semantics/pragmatics distinction as a distinction between what is said (which is semantic) and what is implicated (which is pragmatic).

The “What Is Said” Strategy faces a well-known problem, raised by Stephen Levinson: Grice’s circle. Grice thinks that in order for a hearer to be able to recover what was implicated, she must first know what was said (1989, pp. 30–31). In other words, the process by which a hearer discovers what was implicated requires what was said to be discoverable in advance and available to serve as the input for that process (Levinson 2000, p. 172). But, Levinson contends, the process of

discovering what was said itself depends on “processes that look indistinguishable from implicatures” (2000, p. 186).

For instance, in an utterance of ‘I’ve eaten breakfast,’ what the speaker *says* is typically that she’s eaten breakfast on the day of utterance, and not just that she’s eaten breakfast at all. But the hearer’s discovery of this narrower content seems to rely on something suspiciously similar to the processes used to recover conversational implicatures. In a context in which someone is wondering whether she’s made enough pancakes and I utter ‘I’ve eaten breakfast,’ it would be horribly uncooperative for me to be saying that I’ve eaten breakfast at least once in my life. Knowing that this would be uncooperative, the hearer interprets me as saying I had breakfast *that morning* (Levinson 2000, p. 184). Thus, “what is said seems both to determine and to be determined by implicature” (Levinson 2000, p. 186). A hearer needs to know what was said in an utterance in order to have the necessary input for the pragmatic process of discovering what was implicated, and yet those same pragmatic processes seem needed to discover what was said in the first place.

In my view, this circularity renders the “What Is Said” Strategy unviable. Ultimately, the problem stems from the fact that the notion of what is said is often taken to be based in speakers’ ordinary indirect speech reports, but those ordinary reports seem totally insensitive to the distinction between locutionary acts, and the illocutionary act of assertion (Austin 1962). Thus, because the notion of what is said seems to dance around the semantics/pragmatics line in ordinary usage, it is a poor candidate for distinguishing semantics from pragmatics. In fact, given the inconsistency in our ordinary indirect speech reports, I think it’s a mistake to try to use that notion for any theoretical purposes, unless it is given a clear definition as a technical term (*cf.* Cappelen and Lepore 1997, p. 289; Szabó 2005, pp. 4–5; Camp 2012, p. 623).

Despite the failure of the “What Is Said” Strategy, Grice’s work nicely highlights the importance of the semantics/pragmatics distinction. In order to explain how the professor

communicates that Anne is not qualified by writing ‘Anne speaks fluent English and is always punctual,’ we really do need some independent level of semantic content that can serve as the input for the pragmatic processes that allow the reader to discover what the professor really thinks. Otherwise, there is no way for the process of recovering the implicature to get started. We need a firm distinction between the semantics of an utterance and its pragmatics, where the former is independent from the latter and thus can serve as the input for the hearer’s process of discovering pragmatic content. Distinguishing semantics from pragmatics is essential to explaining how communication works in these kinds of cases.³

Moreover, we can see the need for the distinction in many cases of just speaking literally (*cf.* Bach 2004, p. 29). If you ask me where Steve Jobs was born and I reply with ‘Steve Jobs was born in California,’ I am able to communicate that Steve Jobs was born in California. In other words, on the pragmatic side, that is what my utterance conveys to you. Barring some unusual context, you could not have discovered that pragmatic content unless the composed meaning of the sentence itself provided an independent clue which was able to get your inferential process started. Again, it is clear that we need to distinguish semantics from pragmatics.⁴

³ Although the notion of conversational implicature as a clearly pragmatic phenomenon enjoys fairly widespread acceptance, Ernie Lepore and Matthew Stone have argued that “the category of conversational implicature does no theoretical work” (2015, p. 83). They contend that some phenomena that Griceans have categorized as implicatures are conventional and thus semantic, and that the rest involve mental processes too diverse to be captured by Grice’s Cooperative Principle and maxims (Lepore and Stone 2015, pp. 147–148, 191). Rebutting these points is beyond the scope of the present paper, but I will say that I think it’s a mistake to assume that everything conventional must be semantic. Conventions appear in language at many levels—the level of phonemes, the level of syllable formation, the level of semantics, and even the level of pragmatics. Additionally, Lepore and Stone’s claims about the diversity of interpretive processes center on phenomena such as metaphor, sarcasm, and hinting, for which excellent resources for implicature-based accounts are available.

⁴ We also need to clearly distinguish semantics from pragmatics because theorists often talk past each other due to their different ways of using those terms (Ezcurdia and Stainton 2013, p. xxxi), because the distinction is closely connected to questions about “the propositional content we are responsible for” when we produce utterances (Bianchi 2004, p. 9), and because the distinction plays a prominent role in other philosophical debates such as the debate about contextualism about knowledge (Ezcurdia and Stainton 2013, p. xxx). The distinction can also be seen as more generally “fundamental to philosophical theorizing, because much philosophical theorizing takes the form of claims about the content of philosophically central discourse” (King and Stanley 2005, p. 112).

So, our discussion of Grice has shown that the “What Is Said” Strategy will not work, but the semantics/pragmatics distinction is nonetheless crucial to the project of explaining communication in many cases, both when pragmatic content diverges from semantic content and when it does not. This means that we need to find a way of drawing the distinction that satisfies the following criterion: we must not say of any interpretive process that it both has semantic content as its input, and is used by the hearer to discover semantic content in the first place.

3. Two Context-Based Strategies

Another influential approach to the semantics/pragmatics distinction is to draw the distinction in terms of the involvement of the context of utterance. We’ll discuss two prominent context-based strategies, drawing on work by Jeffrey King and Jason Stanley. This discussion will lead to two additional criteria for a satisfactory treatment of the semantics/pragmatics distinction.

In what I will call the *Strict Context-Based Strategy* (or *Strict Strategy*, for short), no contributions from the context of utterance are permitted into semantics. The semantic content of a complex expression, then, is a matter of just the context-independent meaning of the simple expressions it contains and the compositional rules encoded in the sentence’s syntactic form (King and Stanley 2005, p. 115). For example, if I now utter ‘Today is Thursday,’ the fact that ‘today’ has the context-independent meaning ‘the day of utterance’ is a semantic matter, but the fact that my token of ‘today’ refers to November 2, 2017 is a pragmatic matter, because it relies on a contribution from context. Richard Montague espoused a version of the Strict Strategy (1974, p. 96; qtd. in King and Stanley 2005, p. 116).

Although this strategy may seem less intuitive than the “What Is Said” Strategy, it has the virtue of avoiding circularity. The processes a hearer uses to discover semantic content will involve only knowledge of expression meaning and compositional rules; there is no reason to think that

hearers must already know the semantic content of the entire utterance in order to utilize those processes.

However, the Strict Strategy has a significant downside of its own: it *over-restricts* the domain of semantics. To see how, consider the fact that although the expression ‘today’ changes reference from context to context, the referent of a token of ‘today’ must *always* be the day on which it is uttered. In other words, ‘today’ displays a universal regularity in its reference.⁵ Here it’s helpful to compare ‘today’ with, for instance, ‘January 4, 2011.’ ‘January 4, 2011’ displays a simpler universal regularity of reference—no matter what the speaker intends, tokens of that expression all refer to the same particular day. ‘Today’ differs from ‘January 4, 2011’ in that tokens of it do refer to different days in different contexts, but crucially, *which* day a particular token of ‘today’ refers to is no more up to the speaker in the case of ‘today’ than in the case of ‘January 4, 2011.’ Tokens of both expressions refer to a particular day in a way that is not within the speaker’s control. My suggestion is that the universal regularity of reference of both ‘today’ and ‘January 4, 2011’ implies that their reference is a semantic matter—something the expression itself does, rather than something that a speaker does with her words beyond what the words themselves do.

The *Moderate Context-Based Strategy* (or *Moderate Strategy*), a more popular strategy in contemporary philosophy of language, avoids the Strict Strategy’s defect. According to the Moderate Strategy, the semantic content of a complex utterance is a matter of the context-

⁵ On this point, see John Perry’s discussion of automatic indexicals (2001/2012, p. 68ff.; cf. King and Stanley 2005, p. 113). But, some cases seem to suggest that ‘today’ and other automatic indexicals do not actually exhibit universal regularity of reference. For instance, in answering machine recordings of ‘I am not in the office today,’ ‘today’ seems to refer to the day on which someone hears the recording, rather than the day on which the recording was made (Kaplan 1989a, p. 491 n. 12). However, if we think of answering machines as devices that delay utterances, the event of the recorded message playing when someone calls actually *is* the utterance, and so the token of ‘today’ still refers to the day of utterance (Sidelle 1991; Cohen 2013). There are also what Perry calls “undexical uses” of automatic indexicals: bound variable uses (e.g., ‘Never put off until tomorrow what you can do today’) and anaphoric uses (e.g., ‘Last week I arrived on Wednesday for the reading group. But then I realized the group wasn’t meeting *today* but *tomorrow*’) (2017, pp. 49–50). I’m inclined to treat these undexical uses as non-literal, or as idioms. ‘Never put off until tomorrow what you can do today,’ for instance, is clearly an idiom, and thus its meaning may differ from the composed meanings of its parts. For the anaphora case, we can think of the anaphoric level of content as just pragmatic. It’s fairly intuitive to imagine a hearer interpreting ‘today’ and ‘tomorrow’ as referring to the day of utterance and the day after the utterance, realizing that this is incompatible with the semantics of the rest of the utterance, and then discovering the pragmatic content.

independent meaning of the smaller expressions it contains, the compositional rules encoded in the sentence's syntactic form, and elements of context for which a role is encoded in context-independent meaning. Thus, a contribution to the content of an utterance can be within the domain of semantics even though it depends on the particular context of utterance, but only if a role for context is encoded in the context-independent meaning of the uttered expression (or its parts). Context here includes everything from the time and location of the utterance, to the interlocutors' mental states (King and Stanley 2005, pp. 116–118).

On this view, the fact that 'today' as uttered on November 2, 2017 refers to November 2, 2017 is a semantic matter because a role for context is built into the meaning of 'today.' So, the Moderate Strategy avoids the Strict Strategy's mistake of over-restricting semantics. And yet, if I utter 'Today is Thursday' to someone who knows that I teach an evening course on Thursdays, the fact that my utterance conveys to her that I won't be available for dinner that evening still comes out as pragmatic, because the role context plays in the communication of that further content is not built into the context-independent meaning of the expressions I uttered.

Because the Moderate Strategy allows context to play a role on both the semantic and pragmatic sides, there may seem to be a danger of circularity. The Moderate Strategy does require the hearer to use similar processes on both sides of the divide: to discover both semantic and pragmatic content, a hearer may need to discover the speaker's intentions, as well as pay attention to other aspects of the context. But crucially, the processes the hearer uses to discover semantic content do not rely on preexisting knowledge of semantic content. Instead, they are just lower-level processes of discovering referential intentions and attending to other relevant features of context. There's no reason to think that those processes require semantic content as input.

For instance, if someone utters 'This is delicious' while eating a slice of cake, a proponent of the Moderate Strategy might want to say that the speaker's referential intention when uttering 'this'

plays a role on the semantic side. The hearer can figure out what the speaker likely intends to refer to just by considering which objects are salient and compatible with the semantics of the rest of the utterance. Then, once he's grasped the utterance's semantic content, the hearer might discover another intention of the speaker's: an intention to get the hearer to recognize that the speaker has abruptly changed the topic of conversation to something trivial because of a belief that the hearer's previous utterance was socially inappropriate (*cf.* Grice 1989, p. 35). In this way, the hearer discovers the pragmatic content. So, despite the fact that intention-discovery might play a role in the hearer's recognition of both the semantic and pragmatic content, the hearer does *not* already have to know the semantic content of the entire utterance in order to discover the intention to refer to the cake that is necessary for discovering the semantic content in the first place.

However, the Moderate Strategy does have one significant disadvantage in comparison to the Strict Strategy: the Moderate Strategy makes the semantics/pragmatics distinction much less *sharp*. According to the Strict Strategy, the difference between semantic facts and pragmatic facts is a difference in *kind*—semantic facts have to do only with context-independent meaning and compositional rules, whereas pragmatic facts have to do with elements of the context. According to the Moderate Strategy, on the other hand, some semantic facts are partially determined by elements of the context, just as pragmatic facts are.⁶ The Moderate Strategy faces the project of drawing a line in the field of contextually contributed elements of content (including those that rely on the speaker's mental states) and herding some of them onto the semantic side, others onto the pragmatic side.⁷ It is not hard to see why this might look like a frantic process of “policing the boundary between semantics and pragmatics” (Armstrong and Michaelson 2016, p. 139). If the

⁶ I should note that in this paper, I use ‘determine’ exclusively in the metaphysical sense—that is, in the sense of making some state of affairs obtain, rather than in the sense of *discovering* that some state of affairs obtains.

⁷ King and Stanley, who support a version of the Moderate Strategy, acknowledge that they face this task, at least with respect to speaker intentions. They say that they “see no basis for skepticism about the possibility of distinguishing, in particular cases, those intentions that are semantically relevant from those that are not” (King and Stanley 2005, p. 130). But even if this task could be completed, we would still be without the additional depth that a difference in kind could provide.

difference between semantic and pragmatic phenomena is not a genuine difference in kind, it's easy to understand why the distinction might seem shallow and not terribly significant. A difference in kind would allow for a deeper, more obviously significant distinction.

As a result of our discussion in this section and the previous one, we now have three criteria for a satisfactory treatment of the semantics/pragmatics distinction:

- (1) Avoid circularity: don't say of any interpretive process that it both has semantic content as its input, and is used by the hearer to discover semantic content in the first place.
- (2) Don't over-restrict semantics.
- (3) Make the distinction sharp: construe the distinction as involving a difference in *kind*.

4. The Non-Mental/Mental Strategy

The next step is to provide a strategy for drawing the semantics/pragmatics distinction that satisfies all three criteria. My proposal, the *Non-Mental/Mental Strategy*, is to draw the semantics/pragmatics distinction in terms of the involvement of mental phenomena, rather than in terms of context or what is said. Semantic phenomena, on this view, are entirely non-mental, whereas pragmatic phenomena have to do with speakers' mental states and hearers' inferences about them. The transition I'm suggesting is from the Context-Based Strategies' idea that we need to draw the distinction in terms of how *much* context is allowed into semantics, to an approach in terms of the *kinds* of elements of context allowed into semantics.

Before continuing, I should say a bit more about my notion of a context of utterance, which is quite similar to King and Stanley's. An utterance's context includes all features of its surroundings: the speaker and hearer(s), their mental states, the time and location of the utterance, nearby objects and their relations to each other, recent (or concurrent) utterances and other actions, *etc.* Thus, contexts have both non-mental and mental elements. This notion of a context is

intentionally quite open-ended. It's not possible to specify exactly how far away from the utterance the context extends, but this indeterminacy mirrors the fact that it's not possible to specify in advance exactly how far away from their utterances two interlocutors will go in their use of context.

We'll consider the semantic and pragmatic sides of the Non-Mental/Mental Strategy in turn, beginning with the semantic side. According to the Non-Mental/Mental Strategy, context-independent expression meaning and compositional rules are non-mental and thus can enter in on the semantic side. Non-mental elements of context can also enter in on the semantic side, but only insofar as a role for them is encoded in the context-independent meaning of the uttered expressions. This last qualification is essential because of course we can learn all sorts of information from non-mental elements of context that has nothing to do with language or communication at all, such as when I learn that someone else was previously in the room with my interlocutor by noticing two drinking glasses on the table. It is only when context-independent meaning encodes a role for some non-mental element of context that that element of context factors into semantics.

The Non-Mental/Mental Strategy does not require a particular understanding of what it is for a role for context to be encoded in context-independent meaning. But, whatever the encoding amounts to, it must provide a determinate role for a determinate element of context to play (rather than, perhaps, an open-ended invitation to bring context into the mix). This means that context-independent meaning has to specify exactly *which* elements of context are relevant, and how.

David Kaplan's work provides one way of cashing out the encoding notion that would meet these requirements. On his picture, the context-independent meaning of an indexical expression such as 'I' is a function from contexts to contents (Kaplan 1989a, p. 505). In particular, the function associated with 'I' is a function from contexts to speakers. This clearly provides a determinate role for a determinate element of context to play. However, Kaplan's notion of context is not the broad notion of an utterance's surroundings that we've adopted. On his view, context is just "a package of

whatever parameters are needed to determine ... content” (Kaplan 1989b, p. 591), such as an agent (*i.e.*, speaker), time, and location (Kaplan 1989a, pp. 495, 498, 512). This understanding of context differs from ours, but it is not incompatible with the Non-Mental/Mental Strategy. We could treat Kaplan’s notion of context as a formal representation of some aspects of our broader notion of context, and then adopt the rest of Kaplan’s picture. So, although I want to remain neutral about exactly what the notion of a role for context being encoded in context-independent meaning involves, it’s worthwhile to note that Kaplan’s idea of a function from contexts to contents provides an excellent way to be more specific about it.

Now, for the Non-Mental/Mental Strategy to succeed, everything we’ve placed on the semantic side must actually be non-mental. The idea that non-mental elements of context (such as the time and the location) are non-mental is obviously not controversial, but the same cannot be said for the claim that context-independent expression meaning is non-mental. Context-independent expression meaning is not ordinarily thought of as basic and unanalyzable, and when philosophers provide accounts of the facts in virtue of which expressions have their context-independent meanings—that is, when they offer *foundational theories of meaning*—they almost invariably appeal to mental phenomena, such as intentions and beliefs.⁸ Such foundational theories of meaning imply that semantic facts are partially constituted by mental phenomena. As a result, it’s important to acknowledge that the Non-Mental/Mental Strategy can succeed only if a non-mental foundational theory of meaning can be provided. A non-mental foundational theory of meaning doesn’t have to deny that language users’ *knowledge* of context-independent expression meaning is mental. Their knowledge of meaning just can’t be what constitutes facts about context-independent meaning at the public, group-wide level.

⁸ For instance, consider the foundational theories of meaning offered by Grice (1989), David Lewis (1975), Brian Loar (1976), Paul Horwich (1998, 2005), and Wayne Davis (2003, 2005), among others.

The claim that compositional rules are non-mental requires defense as well, and its defense would involve providing what we might call a *foundational theory of compositional rules*—that is, an account of the facts in virtue of which particular syntactic forms have particular compositional significance. As with the foundational theory of meaning, the facts in virtue of which syntactic forms have their compositional significance must be entirely non-mental in order for the Non-Mental/Mental Strategy to succeed. And again, this claim does not require us to deny the existence of language users' internal competence with compositional rules.

Although the ultimate success of the Non-Mental/Mental Strategy requires the provision of these two non-mental foundational theories, within this paper the project is just to show that the Non-Mental/Mental Strategy lays out a promising way forward for the semantics/pragmatics distinction. The claim that non-mental foundational theories of meaning and compositional rules would be highly desirable can be thought of as a kind of secondary thesis of this paper. I'll say a bit more about this issue at the end. Leaving the issue aside for now, we can summarize the semantic side as follows: a complex utterance's semantic content is determined by the context-independent meaning of the uttered expressions, compositional rules, and non-mental elements of context for which a role is encoded in context-independent meaning.

On the pragmatic side, as I said above, the general idea is that speakers' mental states and hearers' inferences about those mental states are pragmatic. But not *all* of the speaker's mental states play a role in pragmatics, nor do all of the hearer's inferences about the speaker's mental states. We need a principled way of deciding which mental phenomena really do play a role in pragmatics.

The notion of pragmatic content is a helpful starting point. On my view, an utterance's pragmatic content is the content of the mental state(s) that the speaker intends the hearer to infer

her to have, using the utterance as an essential clue.⁹ An inference uses an utterance as an essential clue if it proceeds via a chain of reasoning that could not have reached its conclusion without the utterance playing the role it played. The utterance's intended (and actual) role as an essential clue often (though not always) involves its semantic content.

Mental states of a speaker fall within the domain of pragmatics, then, if they play some role in determining pragmatic content. Using this guideline, mental states that the speaker intends the hearer to discover using the utterance as an essential clue are obviously pragmatic, because it is from them that an utterance inherits its pragmatic content. The speaker's intention that the hearer draw certain inferences about her mental states using the utterance as an essential clue clearly plays a role in determining pragmatic content as well.¹⁰ Such intentions are often complex—that is, the speaker will usually have subsidiary intentions about specific clues the hearer should exploit beyond the utterance itself in order to draw the intended inference about the content of the speaker's mental states. Those clues can include mental states of the speaker's of which the speaker believes the hearer already to be aware, and non-mental elements of the context. Here it's important to emphasize that what can play a role on the pragmatic side is the speaker's *intention* that the hearer use some (possibly non-mental) element of the context as a clue, and not that bare element of context itself.

Whereas the speaker's side of pragmatics has to do with the factors that *determine* pragmatic content, the hearer's side has to do with the process of *discovering* pragmatic content. Specifically, the hearer's inferences about the content of the speaker's mental state(s), using the utterance as an

⁹ This notion of pragmatic content, and in particular, the role for the speaker's intention about what the hearer should infer, is influenced by (though certainly not identical to) Grice's (1989) notion of speaker meaning.

¹⁰ In cases in which a speaker is attempting to deceive a hearer, there will of course be no actual mental state of the speaker's of which it is true that she intends the hearer to discover its content. Rather, there will be only the speaker's intention that the hearer draw certain inferences about the content of the speaker's mental states.

essential clue, enter in on the pragmatic side. This is the moment at which communication succeeds, if it succeeds at all.

Without the qualification that the hearer's inferences must rely on the utterance as an essential clue in order to play a role in pragmatics, we would be allowing into pragmatics inferences of the hearer that are just detective work about mental states, rather than a part of the communication process. For instance, if I notice a suitcase by my friend's door and infer that he intends to go on a trip, my inference about the content of his mental state has nothing to do with communication. On the other hand, if my friend utters 'I'm looking forward to getting out of here for a while,' I may discover the same intention, but this time in an utterance-dependent manner that makes it a case of communication rather than of mere detective work about mental states.

We are now in a position to fully describe the Non-Mental/Mental Strategy. Semantic phenomena are a matter of context-independent meaning, compositional rules, and the contributions of non-mental elements of the context of utterance for which a role is encoded in context-independent meaning. Pragmatic phenomena can be divided into two kinds. First, the following mental states of the speaker enter in on the pragmatic side: mental states the speaker intends the hearer to discover by using the utterance as an essential clue, the speaker's intentions about what inferences the hearer should draw about her mental states using the utterance as an essential clue, and the speaker's subsidiary intentions about what clues beyond the utterance the hearer should exploit in making those inferences. Second, the hearer's inferences about the content of the speaker's mental state(s), using the utterance as an essential clue, enter in on the pragmatic side as well.

Let's apply the Non-Mental/Mental Strategy to some of our examples. When the professor writes 'Anne speaks fluent English and is always punctual' in a letter of recommendation, the semantic content, relying on context-independent meaning, compositional rules, and the time at

which the utterance occurred, is a straightforward attribution of two good qualities to Anne. But then, the professor intends the reader to see that in order to preserve the assumption that the professor is being cooperative while attributing such trivial virtues to Anne, he must infer that the professor believes something else. This process relies essentially on the clue provided by the utterance, but the professor also intends the reader to notice that the utterance occurs in the context of a letter of recommendation. The pragmatic content, then, is that Anne is not qualified for graduate school—this is the content of the professor’s belief, as intended to be inferred by the hearer.

On the other hand, when I utter ‘Today is Thursday,’ the semantic content is that November 2, 2017 is a Thursday. This comes from the context-independent meaning of the words and the relevant compositional rules, along with one non-mental element of the utterance’s context: the day on which the utterance occurs. And, assuming that my utterance is addressed to someone who knows I teach a course on Thursday evenings, the pragmatic content could be that I will be unable to have dinner with the hearer that evening. My intention that she make use of her memories of what I have done on past Thursdays also factors in on the pragmatic side, in a subsidiary capacity.¹¹

And, to consider the simplest kind of case, when you ask where Steve Jobs was born and I reply by uttering ‘Steve Jobs was born in California,’ the semantic content is that Steve Jobs was born in California. I intend for you to grasp this semantic content and, noticing nothing anomalous about it as a contribution to the conversation, infer that I believe that Steve Jobs was born in California. The content of that belief, then, is the pragmatic content of my utterance. In this case, the pragmatic content matches the semantic content.

¹¹ Strictly speaking, the pragmatic content in both examples is two-fold. In the first case, the professor communicates that Anne speaks fluent English and is always punctual, and also that Anne is not qualified for graduate school. In the second, I communicate that November 2, 2017 is a Thursday, as well as that I will be unable to have dinner with the hearer that evening. I confined my attention to the more interesting part of the pragmatic content in the main text, for ease of presentation.

Now we'll apply the criteria from the end of Section 3 to the Non-Mental/Mental Strategy. In satisfaction of our first criterion, there is no circularity in the Non-Mental/Mental Strategy's way of drawing the distinction. It's true that in order to discover pragmatic content, the hearer often relies on the semantic content as an essential clue. But crucially, to discover semantic content, the hearer requires only knowledge of expression meaning, compositional rules, and non-mental elements of context. The hearer does *not* utilize any processes that require preexisting knowledge of the utterance's semantic content.

The Non-Mental/Mental Strategy also avoids the Strict Strategy's mistake of over-restricting semantics, in connection to our second criterion. To return to the 'Today is Thursday' example, on the semantic side, we still have context-independent meaning and compositional rules, but the day on which the utterance occurs *does* play a role on that side as well, because it is a non-mental element of the context of utterance for which a role is encoded in the context-independent meaning of 'today.' In Section 5, however, we'll see that there's more to be said about the Non-Mental/Mental Strategy in relation to semantic over-restriction.

Additionally, the Non-Mental/Mental Strategy satisfies our third criterion by making the semantics/pragmatics distinction *sharp*. Just like on the Strict Strategy, we really do have a difference in kind here—non-mental versus mental phenomena. We're not trying to draw a line in the middle of an open field of elements of context (or anything else). Instead, we've identified a deeper difference between semantic and pragmatic phenomena.

In fact, the difference in kind we've identified goes one level deeper than on the Strict Strategy. On the Strict Strategy, the pragmatic side consists of a unified kind: contributions from context. But the Strict Strategy does not identify a unified kind to which all semantic phenomena belong. Instead, we have two kinds of things on the semantic side: context-independent expression meaning, and compositional rules. Now of course, all of the phenomena on the Strict Strategy's

semantic side *are* among the phenomena I've been characterizing as non-mental. However, the Strict Strategy also categorizes some non-mental phenomena—namely, non-mental elements of context—as pragmatic. By providing a unified description of the kind of phenomena that appear on *each* side of the divide, the Non-Mental/Mental Strategy grounds the semantics/pragmatics distinction in a deeper difference in kind than even the Strict Strategy provided.

At this point, a reader with strong materialist sensibilities about the mind may be inclined to challenge the claim that the non-mental/mental distinction tracks a genuine difference in kind. But even if we grant that the difference between mental and non-mental phenomena isn't metaphysical, non-mental and mental phenomena still differ in kind in virtue of their *epistemic* properties. Non-mental phenomena are *epistemically public* in the sense that all minds have equal access to them, *ceteris paribus*. For instance, any of us could discover facts about the sizes of distant planets, if we were in the right place, with the right tools. The same goes for facts about context-independent expression meaning and compositional rules—anyone can access these facts so long as they have encountered tokens of the relevant expression or applications of the rule, and assuming they have the right tools, such as knowledge of the language or a translation dictionary.

On the other hand, even under materialist assumptions, mental phenomena are *epistemically private* in the sense that a single mind has privileged access to them, *ceteris paribus*. Your mental states right now are epistemically private because you have a kind of epistemic access to them (and to facts about them) that others lack, and which they could not gain by being “in the right place with the right tools.” Even if scientists eventually developed a way of reading the content of mental states directly off of brain scans, they could not achieve the kind of immediate access to your mental states that you enjoy. We can thus think of non-mental and mental phenomena as different in kind as a result of this difference in epistemic access to them.

5. Demonstratives and Semantic Over-Restriction

In this section and the two that follow, we'll consider three different challenges to the Non-Mental/Mental Strategy, beginning with a new worry about semantic over-restriction.

Imagine that someone utters 'That is beautiful' while standing in an art museum, surrounded by many works of art. It is natural to think that at least part of what determines the referent of the speaker's token of 'that' is her referential intention—that is, her intention to pick out one particular work of art (see, *e.g.*, Kripke 1977; Kaplan 1989b; Bach 1992; Perry 2001/2012, 2009; Åkerman 2009). But at the same time, although bare demonstratives do not exhibit the obvious universal regularity of reference that we noted in the case of 'today,' many people have thought the reference of an uttered demonstrative to be a semantic matter (*e.g.*, Kripke 1977; Kaplan 1989a; Reimer 1991; Borg 2004, 2012; Cappelen and Lepore 2005; King 2013; Speaks 2016).

If it were true that the referent of an uttered demonstrative is determined by the speaker's intentions and that it is a semantic matter, this fact would be incompatible with the Non-Mental/Mental Strategy's exclusion of all mental phenomena from semantics. We would have a role for a *mental* element of the context encoded in context-independent meaning, which the Non-Mental/Mental Strategy cannot countenance. This would make the Non-Mental/Mental Strategy guilty of a new form of semantic over-restriction.

One possible response would be to argue that the reference of an uttered demonstrative is not a semantic matter, but I'm not inclined toward that approach. Rather, I want to provide an outline of an account of demonstratives according to which the reference of a token demonstrative does *not* depend on the speaker's intentions (or on any other mental phenomena). This will suggest that the reference of an uttered demonstrative can be classified as semantic without undermining the Non-Mental/Mental Strategy.

Here I'll rely on some work by Christopher Gauker (2008), which builds on earlier work by Howard Wettstein (1984). Gauker offers a semantic account of bare demonstratives that excludes speakers' intentions. Instead, he presents a set of criteria that he thinks work together to determine the reference of a token demonstrative. Because Gauker does not share my aim of making semantic content a matter of only non-mental elements of context, mental states (aside from intentions) seem likely to play a role in some of his criteria. Below is an adaptation of Gauker's criteria designed to exclude all mental phenomena, not just intentions:

- (1) Causal access to the hearer's senses. The object stands in a relation to the hearer such that it can have an impact on her sense organs. In other words, it's something the hearer can perceive. This relation can be mediated by devices such as televisions.
- (2) Prior reference. The object was referred to earlier in the conversation. This prior reference must be semantic, not merely pragmatic.
- (3) Relevance. The object is closely connected to recent actions or to the semantic features of recent utterances. This is a matter of relations such as causation, similarity, or shared origin, and not of the interlocutors' mental states.
- (4) Compatibility with the semantics of the rest of the utterance. For instance, in a room full of works of art, the trashcan in the corner is not a good candidate for being the referent of the token of 'that' in an utterance of 'That is beautiful.'
- (5) Pointing. The object intersects (or nearly intersects) a line extending from the speaker's pointing gesture.
- (6) Location in a series. If the speaker and hearer are in an environment in which objects are presented in a series, being next in the series supports an object's candidacy for reference. Here Gauker gives an example of two quality control workers at a factory who utter 'That's good' or 'That's bad' as items come down the assembly line (2008, pp. 364–365).

Crucially, the satisfaction of each criterion comes in degrees. For instance, an object to which someone referred thirty seconds ago satisfies criterion (2) *better* than one to which someone referred two minutes ago. Multiple objects may satisfy multiple criteria to some degree, but the one that best satisfies the overall set of criteria is the referent. The criteria all have equal weight at the general level, though a greater degree of satisfaction of one criterion on some occasion will cause that criterion to have a greater degree of influence on the determination of reference on that occasion (*cf.* Gauker 2008, pp. 364–366).¹²

To see how the criteria work together to determine reference, consider the ‘That is beautiful’ example again, with the following additional details: there is no pointing gesture (so criterion (5) doesn’t come into play), the speaker and hearer are standing in front of Van Gogh’s *The Starry Night* with their faces turned toward it, and the most recent previous utterance in the conversation (which occurred several minutes prior to the utterance of ‘That is beautiful’) semantically referred to a sculpture in another room of the museum. In this scenario, all of the works of art in the museum do pretty well with respect to the relevance criterion, criterion (3), because the speaker and hearer have been viewing and talking about art located inside the museum. Many of the works of art in the museum are also compatible with the semantics of ‘is beautiful,’ in relation to criterion (4). Criteria (3) and (4), then, provide a wide range of possible referents, including both the sculpture and *The Starry Night*. Criterion (2) clearly favors the sculpture, because the most recent prior semantic reference was to it. On the other hand, even though several of the works of art favored by criteria (3) and (4) probably satisfy criterion (1) to some degree, criterion (1) favors *The Starry Night* by a

¹² Two differences between Gauker’s view and the adapted version just outlined are worth mentioning. First, Gauker thinks that our *judgments* about which object best satisfies the criteria are what actually determine reference (2008, p. 369). This is not something I can endorse because it gives mental phenomena (namely, judgments) a role to play in determining the semantic features of utterances that include token demonstratives. On my version of the view, the criteria directly determine reference. Second, Gauker requires that the referent “adequately satisf[y] the criteria,” rather than just requiring (as I do) that the referent satisfy them *best* (2008, p. 364). I don’t think this notion of adequacy is necessary, and furthermore, without a connection to our judgments of adequacy, it seems impossible to specify a standard for adequate satisfaction of the criteria.

wide margin because the fact that the hearer is facing *The Starry Night* means that it can have a much more significant impact on her sense organs than the other works of art can. Because the degree to which *The Starry Night* satisfies criterion (1) is greater than the degree to which the sculpture satisfies criterion (2), and because the sculpture and *The Starry Night* are tied with respect to criteria (3) and (4), *The Starry Night* is the referent of the token of ‘that.’ In fact, criterion (6) favors *The Starry Night* as well, since working one’s way through an art museum can be thought of as a way of having objects presented in a series.

Given that we’re working with six criteria, it’s easy to imagine many contexts in which things don’t go as smoothly as in the *Starry Night* case. It seems likely that in many possible contexts, one object could be most relevant, another object could have the most access to the hearer’s sense organs, and yet another could be most compatible with the semantics of the rest of the utterance. In such cases, the six criteria would fail to identify one object as the referent (*cf.* Gauker 2008, p. 366). Thus, the account of demonstratives we’ve outlined implies that bare demonstratives have significant potential for referential failure.

This implication might seem to be a mark against the proposed account of demonstratives. After all, we do not seem to experience frequent referential failure in our everyday use of bare demonstratives. However, demonstratives’ vulnerability to referential failure in many easily imagined possible contexts does not necessarily imply that demonstratives actually *do* often fail to refer. It is part of our skill as language users that we typically don’t use bare demonstratives unless the context is one in which there is a clear winner, which means that most cases in which bare demonstratives are actually used are like the *Starry Night* case. In contexts that are unfriendly to bare demonstratives, we tend to use other referring expressions such as complex demonstratives, proper names, or definite descriptions. For instance, if in the *Starry Night* case the prior reference to the sculpture had been within the past few seconds, the speaker might have uttered ‘That painting is

beautiful' instead of just 'That is beautiful.' A bare demonstrative would not have been effective in that context, and a skillful speaker wouldn't use one.

Of course, we are not infallible, and we do sometimes use bare demonstratives in contexts in which no single object satisfies the criteria better than all others. In such cases, the criteria really do imply that the token demonstrative fails to refer, though the speaker might still succeed in pragmatically conveying information to the hearer about some object the speaker has in mind (perhaps because of the hearer's pre-existing knowledge of the speaker's mental states). But I see nothing problematic about an account of demonstratives that implies that referential failure occurs on a limited scale, when language users make mistakes.

Coming from another direction, one might object that although the criteria seem to generate the right results in the *Starry Night* case and others, that is only because they are the clues used to discover the object to which the speaker intends to refer, and the speaker's *intention* is nonetheless what actually determines reference. To assuage this worry, it will be useful to consider a case in which the verdict of the criteria diverges from the speaker's intentions.¹³

Imagine that you and I are driving down a highway together, passing the time by reminiscing about a concert we attended on April 11, 2016. We then discuss another happy day—July 16, 2016, when we attended a mutual friend's wedding. The conversation pauses, and my thoughts drift back to the events of April 11. I forget that we ever discussed the July 16 wedding. I utter 'That was a great day,' intending to refer to April 11, when we saw the concert. April 11 and July 16 seem to be tied with respect to all criteria aside from criterion (2): although both days satisfy criterion (2) to some degree, July 16 satisfies it better because the prior reference to it was more recent. So, in this case, the object that best satisfies the criteria is not the same as the object to which I intend to refer. This strikes me as a case in which the *speaker's referent* (on the pragmatic side) is April 11, but the

¹³ Gauker (2008, p. 363) and Wettstein (1984, pp. 70–71) also provide examples of this sort.

semantic referent is July 16. This suggests, then, that the criteria (or something similar to them), and not the speaker's intentions, determine the reference of token demonstratives.¹⁴

My aim in this section has not been to definitively argue for this particular approach to demonstratives, but rather just to show that the Non-Mental/Mental Strategy has resources for classifying the referent of an uttered demonstrative as semantic, thereby avoiding the threat of semantic over-restriction. To give a complete and convincing account of demonstratives, more detail would have to be provided about how the criteria are balanced with respect to one another in a wider range of cases, and the possibility of adding other criteria to the list would need to be discussed. But still, it's interesting to note that if a criteria-based account of demonstratives along these lines is correct, demonstratives actually do exhibit a kind of universal regularity of reference, despite appearances to the contrary. A token demonstrative does always refer to the same element of context—the one that best satisfies the criteria. This is much more complex than indexicals' universal regularities of reference, but it is universal regularity nonetheless.

Now of course, bare demonstratives are not the only kind of expression whose context-independent meaning might reasonably be thought to include a role for mental states, thereby posing a similar threat to the Non-Mental/Mental Strategy. Others include complex demonstratives, pronouns, and indefinite descriptions. I focused on bare demonstratives because they strike me as

¹⁴The 'speaker's referent' and 'semantic referent' terms come from Saul Kripke (1977). In connection to Kripke's work, it's worthwhile to discuss an additional worry about the account of demonstratives we've outlined. Imagine that, before a meeting begins in the conference room at A and B's workplace, A utters 'James is back on the throne.' The speaker's referent of A's token of 'the throne' is the chair at the head of the conference table, where James insists on sitting for all meetings. Semantically, that token of 'the throne' fails to refer. Now suppose the chair is in rather poor condition. In reply, B utters 'I don't envy him. I wouldn't want to sit on that.' It might seem that B's token of 'that' semantically refers to the chair at the head of the table because of A's prior reference to that chair, despite the fact that the chair was not the semantic referent of A's token of 'the throne,' which goes against my statement of criterion (2). This case is analogous to cases Kripke discusses involving pronouns that appear to take as their semantic referent the object that was merely the speaker's referent in a previous utterance (2013, pp. 127–129). However, I think we can explain why the chair is the semantic referent of B's token of 'that' using criteria other than prior reference. Although there was no prior semantic reference to the chair to which the token of 'that' refers, the chair does bear an important connection to the semantic features of the token of 'throne' that was uttered—the chair is not a throne, but it is something that James treats as if it were a throne, which makes that chair satisfy criterion (3) better than the other chairs in the conference room. That chair also satisfies criterion (4) better than the others in the room because it is in poor condition and thus less suited to sitting. A similar solution is, I believe, available for Kripke's cases involving pronouns.

the most difficult case—the case in which context-independent meaning seems to provide the least amount of constraint, and where the need for the speaker’s intentions thus seems greatest. My hope is that taking on the particularly tricky case of bare demonstratives has provided reason for optimism that the same could be done for other kinds of expressions as well.¹⁵ But ultimately, in order for the Non-Mental/Mental Strategy to succeed, non-mental treatments of all semantically context-sensitive expressions are needed.

6. Scalar Implicature and Pragmatic Over-Restriction

From another direction, one might also worry that the Non-Mental/Mental Strategy over-restricts pragmatics. There are certain communicated contents that we ordinarily think of as merely pragmatic, but which hearers seem to discover without thinking about speakers’ mental states. For instance, consider scalar implicature. If one teacher addresses an utterance of ‘Some of the children are lost’ to another during a field trip, she seems to implicate that not *all* of the children are lost. Because these implicatures are so common, hearers tend to grasp them immediately, without seeming to think about the speaker’s mental states. This then suggests that in a case of scalar implicature, the speaker—who, after all, has been a hearer in the past—may also not intend the hearer to think about the speaker’s mental states while discovering the implicated content. Thus, it seems that the Non-Mental/Mental Strategy cannot classify the content of a scalar implicature as pragmatic.¹⁶

The reason this outcome is problematic, of course, is that there’s strong reason to think that scalar implicatures are pragmatic: they are cancelable (Grice 1989, p. 39). That is, the teacher can go

¹⁵ There is also promising work available concerning pronouns. Una Stojnic, Matthew Stone, and Ernie Lepore argue that the semantic value of a token pronoun is determined by which entity is “at the ‘center of attention’” on that occasion. Although the ‘attention’ terminology sounds likely to bring in mental states, in fact it does not: “what lies at the center of attention is *not* governed by ... speaker intentions ..., but entirely by linguistic rules” (Stojnic *et al.* 2017, p. 521).

¹⁶ I’m grateful to an anonymous referee for bringing this objection to my attention.

on to utter ‘Moreover, *all* of them are missing,’ without inconsistency. If the teacher’s initial utterance semantically implied that not all of the children were lost, she would have to contradict herself or retract her first utterance if she later wanted to affirm that all of them were gone.

The idea of cancelability as an indication that some phenomenon is not semantic is, at least in general, compatible with the Non-Mental/Mental Strategy: we can say that the cancelability of scalar implicatures shows that the implicated content must not have come from the composed meaning of the sentence or contributions from context for which a role is encoded in context-independent meaning. Otherwise, the teacher would not be able to follow ‘Some of the children are lost’ with ‘Moreover, *all* of them are missing’ without contradicting herself.

Taken together, the above considerations suggest that the Non-Mental/Mental Strategy leaves scalar implicatures in a funny middle ground—not semantic, because cancelable, but not pragmatic either. However, I’m not convinced that scalar implicatures don’t involve actual and intended inferences about the content of speakers’ mental states. Although hearers don’t have to think consciously and explicitly about speakers’ mental states when they recover scalar implicatures, we can think of them as relying on a defeasible *assumption* about speakers’ mental states. Because scalar implicatures are so common, we have a standing assumption about the kinds of mental states speakers have when they utter sentences that could generate such implicatures. Similarly, because we know we’ll be interpreted in accordance with such assumptions when we utter sentences that could generate scalar implicatures, we typically utter such sentences only with the intention to be interpreted in that way. Thus, the Non-Mental/Mental Strategy can classify scalar implicatures as pragmatic, avoiding pragmatic over-restriction.

As evidence for the plausibility of this view, consider what would happen if we somehow stopped uttering the constructions that typically generate scalar implicatures. It’s easy to imagine scalar implicature being reintroduced by a creative speaker, with the hearer discovering the content

of the speaker's mental state by relying on Grice's first maxim of quantity (1989, p. 26; *cf.* Millikan 2005, pp. 31–32). This suggests that our ordinary process of recovering scalar implicatures can be thought of as a shortcut version of an inference appealing to the first maxim of quantity. Instead of working out each scalar implicature by explicitly appealing to the maxim, we save time by immediately inferring the scalar implicature whenever we encounter utterances of the right form.

Analogously to our discussion of semantic over-restriction in Section 5, we've just considered *one* kind of use of language that seems to suggest that the Non-Mental/Mental Strategy leads to pragmatic over-restriction. In reality, there are many instances in which people discover content that we'd like to categorize as pragmatic, without going through an explicit, conscious process of reasoning about the speaker's mental states. But this discussion of scalar implicature can serve as the model for how to handle other cases of that sort.

7. Intentions Prior to Semantics

Now we'll turn to a final worry about the Non-Mental/Mental Strategy. The strategy aims to ensure the sharpness of the semantics/pragmatics distinction by keeping mental phenomena entirely out of semantics. But there is reason to think that something mental—specifically, the speaker's intentions—must sometimes play a role *prior* to semantics. This occurs in connection to the issue of determining which expression a speaker uttered, which thereby partially determines the utterance's semantics. If mental states were necessary at that pre-semantic level, then the Non-Mental/Mental Strategy would be unviable.

Kepa Korta and John Perry provide a case that they believe illustrates the role of intentions in determining the language to which some utterances belong, which thereby partially determines which semantic features such utterances have. In their example, John arrives in the Basque Country and greets his two hosts with an utterance that is acoustically indistinguishable between the English

sentence ‘Nina is John’ and the Basque sentence ‘Ni naiz John.’ The latter is a word-by-word translation of ‘I am John’ into Basque. One of the hosts thinks John uttered ‘Nina is John,’ perhaps as the set-up for a joke; the other thinks he uttered ‘Ni naiz John,’ attempting to introduce himself in Basque. Korta and Perry contend that when we ask which sentence John uttered, “[t]he answer seems to be provided in large part by what John was *trying* to do, what his intentions were in making the sounds he did” (2011, pp. 1–2). Because the sound John produced equally resembles ‘Nina is John’ and ‘Ni naiz John,’ John’s intentions seem needed to determine which expression he uttered.

If this were true, it would be a significant problem for the Non-Mental/Mental Strategy. If speakers’ intentions played a role in determining the language to which some utterances belong, those intentions would thereby also play a role in determining which semantic features those utterances have. On this picture, if John intended to utter a token of the English sentence, the first two syllables of the sound he made would have the semantics of ‘Nina,’ thereby referring to the bearer of that name. On the other hand, if he intended to utter a token of the Basque sentence, the first syllable of his utterance would have the same semantics as ‘I’ does in English. If speakers’ intentions sometimes had to come in to help determine which semantic features an utterance has, then some semantic facts would be partially determined by mental phenomena. This would undermine the Non-Mental/Mental Strategy—we would no longer be able to say that mental phenomena play no role on the semantic side.

For our purposes, there is nothing special about the fact that the two candidate sentences belong to different languages in Korta and Perry’s case. Disambiguation within a single language gives rise to the same issue. If I utter ‘The bark was impressive,’ it may seem that we need to appeal to my intentions to discover whether the semantic content of my utterance concerns a dog’s cry or the surface of a tree trunk. Thus, the issue again is that intentions seem needed to play a role in

determining which expression a speaker's utterance tokens, thereby undermining the Non-Mental/Mental Strategy.

My contention is that, despite appearances, we are *not* forced to appeal to the speaker's intentions. There is another, non-mental option for explaining what determines the type to which an uttered token belongs. Ruth Garrett Millikan argues that the type to which a token expression belongs is determined entirely by its causal history (2005, pp. 33–35). Thus, whether the semantics of my utterance concerns a dog's cry or the surface of a tree trunk is determined by the chain of past uses that influenced my current one. Am I influenced by the chain of past uses directed toward sounds made by dogs, or toward the surfaces of tree trunks? *That*, along with compositional rules and some non-mental elements of context, is what settles the semantics of my utterance—not my intentions.¹⁷ Similarly, we can say that whether John's utterance had the semantics of 'Nina is John' or of 'Ni naiz John' is a matter of the chain of past uses that influenced him, not a matter of his intentions. If he was influenced by chains of past utterances of 'Nina' and 'is,' he uttered 'Nina is John.' If he was influenced by chains of past utterances of 'ni' and 'naiz,' he uttered 'Ni naiz John.'¹⁸

Importantly, the causal chain leading up to an utterance is non-mental. Unlike my intention to say something about the cry of a dog or John's intention to token English or Basque expressions, the connection between an utterance and a long history of past utterances is not a state of anyone's mind; it is a chain of causal relationships among actual events of linguistic behavior. So, on this

¹⁷ On Millikan's picture, there are two kinds of cases: cases in which there are two entirely independent chains of past usage that just happen to involve acoustically similar sounds (as in the case of 'bark'), and cases in which there are what Millikan calls "branching" chains of past usage, which give rise to different, but often closely related, senses of a single sound (such as the use of 'lip' to mean the flesh surrounding a mouth, versus its use to mean the rim of a glass) (2005, pp. 33–35, 61).

¹⁸ I don't mean to obscure the fact that some utterances are influenced by multiple chains of past utterances. If one of those chains is the *primary* influence, that chain determines which expression was tokened, and thereby, the utterance's semantics. For instance, if I produce a sound that is primarily influenced by chains of past uses of 'espresso' but also is influenced enough by past utterances of 'express' that the pronunciation of the first 's' sounds more similar to an 'x,' the dominant causal influence prevails, and my utterance is of 'espresso,' improperly pronounced. But cases in which multiple causal chains exert equal influence are at least possible, and in such cases, which expression the speaker tokened (and thus, the utterance's semantics) is indeterminate.

picture, no mental phenomena have to come in prior to semantics to determine an utterance's semantic features.

Now of course, typically a speaker's intentions when producing an utterance are in agreement with the causal history that influenced the utterance. What's important for our purposes is that we *need not* appeal to those intentions in order to explain the fact that an utterance, despite its resemblance to two different expressions, has one determinate semantic content. There is a viable non-mental explanation available, in terms of the causal history of the utterance. Thus, there is no threat to the Non-Mental/Mental Strategy.¹⁹

In Section 6, I based my argument against the threat of pragmatic over-restriction largely on the availability of a mental explanation of how the content of a scalar implicature is conveyed, but without ruling out possible non-mental explanations. In the present section, I have based my argument against the pre-semantic intrusion of mental phenomena on the availability of a purely non-mental explanation of how the type of an utterance is determined, but without ruling out possible mental explanations (such as the one Korta and Perry offer). In Section 5, too, significant weight was placed on the availability of a non-mental explanation of how token demonstratives refer. Taking all three sections together, it may seem that I'm cherry-picking—choosing non-mental or mental explanations when it suits me, rather than showing that any of these phenomena actually are non-mental or mental.

¹⁹ Having brought Korta and Perry's work into the mix, I should also say something about their treatment of the semantics/pragmatics distinction. Their characterization of pragmatics is similar to mine: they see it as closely connected to speakers' intentions. But their construal of the domain of semantics implies that in an utterance of 'That is beautiful,' the token of 'that' would not have a referent at the level of semantic content (Korta and Perry 2011, pp. 54, 140; *cf.* 2008, p. 351). Thus, they over-restrict semantics in the way we took pains to avoid in Section 5. Kent Bach's conception of pragmatics is even more similar to mine in that he thinks pragmatic content is the content of the mental states a speaker intends the hearer to infer her to have (2002, p. 286). But like Korta and Perry, he over-restricts semantics. He thinks that demonstratives do not have semantic referents, and that as a result the output of semantics for a sentence containing a demonstrative is not a complete proposition (Bach 2004, pp. 36–37). Another difference between my view and Bach's is that Bach thinks only sentences have semantic features, whereas I think utterances, too, have semantic (as well as pragmatic) features (2004, p. 28).

However, I do think that the mere availability of explanations compatible with the Non-Mental/Mental Strategy is significant, even if explanations incompatible with the strategy are available as well. The fact that the strategy points the way toward a sharp, non-circular semantics/pragmatics distinction, combined with the semantics/pragmatics distinction's crucial role in the project of explaining communication (as discussed in Section 2), itself provides a reason to prefer accounts of various phenomena that are compatible with the strategy. Thus, though it would clearly be a problem for the Non-Mental/Mental Strategy if explanations compatible with it *weren't* available, we don't need to do more than establish that such explanations are in fact available in order to provide support for a strategy that has already been shown to be desirable on other grounds.

8. Conclusion

We've addressed three challenges to the Non-Mental/Mental Strategy: we've shown how the semantic context-sensitivity of demonstratives could be accommodated without appealing to mental phenomena, we've discussed the strategy's resources for categorizing scalar implicature as pragmatic, and we've addressed the worry that mental phenomena might sometimes play a role prior to semantics. I'd like to end by outlining the projects that need to be undertaken in order to move from the Non-Mental/Mental Strategy—which is, after all, just a strategy—to an actual *account* of the semantics/pragmatics distinction in non-mental/mental terms.

First, as we discussed in Section 4, the Non-Mental/Mental Strategy ultimately cannot succeed without entirely non-mental foundational theories of both expression meaning and compositional rules. Thus, providing those two foundational theories is a crucial next step toward a non-mental/mental account of the semantics/pragmatics distinction.²⁰ My view is that this can be

²⁰ One could avoid this step by treating the context-independent meaning of expressions and the compositional significance of syntactic forms as basic. This would amount to declining to give foundational theories at all, which would certainly prevent mental phenomena from sneaking in. Nathan Salmon takes this approach with respect to

done by focusing on linguistic behavior. Our linguistic behavior—our production of various sounds, gestures, and markings in various sequences—is non-mental. If we could explain the context-independent meaningfulness of linguistic expressions and the compositional significance of syntactic forms as arising solely from behavior, we would have the foundational theories we’re looking for. But that is a project for another time.²¹

Second, as we discussed at the end of Section 5, we need to show that utterances of all semantically context-sensitive expressions can be assigned their semantic contents without appeal to intentions or other mental states. I hope the discussion of the particularly challenging case of bare demonstratives has gone some of the way toward showing how this could be done. At the end of that section, I acknowledged other kinds of referring expressions that need to be considered, but of course other kinds of context-sensitivity, such as that involved in the tense of verbs, ultimately need to be considered as well. Additional work related to avoiding pragmatic over-restriction would be helpful, too, although for pragmatic over-restriction, the solution for further cases will resemble the solution in the case of scalar implicature much more closely than accounts of different semantically context-sensitive expressions will resemble each other.

It’s clear that much work remains to be done to move from our strategy for drawing the semantics/pragmatics distinction to a robust account of the semantics/pragmatics distinction. But for the present, I will be satisfied if this discussion of the Non-Mental/Mental Strategy has provided

context-independent meaning: “the semantic attributes of expressions are ... intrinsic to the expressions themselves, or to the expressions *as* expressions of a particular language” (2005, p. 324). I won’t say much about this possibility because expression meaning and compositional significance strike me as highly unlikely to be basic or intrinsic. Or, if we emphasize Salmon’s qualification that semantic features are intrinsic to expressions only “*as* expressions of a particular language,” then the interesting question becomes: What makes an expression (at the type level) belong to a particular language? The same threat of the intrusion of mental states—and thus to the Non-Mental/Mental Strategy—would then reappear at that level. An analogous point can be made about the option of treating the compositional significance of syntactic forms as basic and intrinsic.

²¹ Michael Johnson and Jennifer Nado have recently offered a foundational theory of meaning that they describe as behavioral. However, they ultimately ground the meaningfulness of linguistic expressions in *dispositions* toward linguistic behavior, and dispositions toward behavior are surely grounded in mental phenomena (Johnson and Nado 2014, pp. 81–82). Brian Skyrms’s work on meaning is also quite behavioral in its focus, but his project is to identify causal mechanisms by which expression meaning can emerge, rather than to give a foundational theory of meaning or of compositional rules (2010, p. 1, Ch. 12). Thus, the need for behavioral foundational theories remains.

reason for optimism that a sharp semantics/pragmatics distinction is possible, and that the terms ‘semantic’ and ‘pragmatic’ still deserve a place in our philosophical lexicon.

Acknowledgements

I’m grateful to audiences at Vassar College, McMaster University, the 2017 Western Canadian Philosophical Association meeting, and the Fall 2017 New York Philosophy of Language Workshop for discussion of past versions of this paper. I’m also grateful to Bradley Shubert for research assistant work, and to two anonymous referees for *Synthese* for extremely helpful feedback.

References

- Åkerman, J. (2009). A plea for pragmatics. *Synthese*, 170(1), 155-167.
- Armstrong, J., & Michaelson, E. (2016). Introduction for *Inquiry* symposium on *Imagination and convention*. *Inquiry*, 59(2), 139-144.
- Austin, J.L. (1962). *How to do things with words* (2nd ed.). Cambridge: Harvard University Press.
- Bach, K. (1992). Intentions and demonstrations. *Analysis*, 52(3), 140-146.
- Bach, K. (2002). Semantic, pragmatic. In J. K. Campbell, M. O’Rourke, & D. Shier (Eds.), *Meaning and truth* (pp. 284-292). New York: Seven Bridges Press.
- Bach, K. (2004). Minding the gap. In C. Bianchi (Ed.), *The semantics/pragmatics distinction* (pp. 27-44). Stanford: CSLI Publications.
- Bianchi, C. (2004). Semantics and pragmatics: The distinction reloaded. In C. Bianchi (Ed.), *The semantics/pragmatics distinction* (pp. 1-11). Stanford: CSLI Publications.
- Borg, E. (2004). *Minimal semantics*. New York: Oxford University Press.
- Borg, E. (2012). *Pursuing meaning*. Oxford: Oxford University Press.
- Camp, E. (2012). Sarcasm, pretense, and the semantics/pragmatics distinction. *Noûs*, 46(4), 587-634.
- Cappelen, H. (2007). Semantics and pragmatics: Some central issues. In G. Preyer & G. Peter (Eds.), *Context-sensitivity and semantic minimalism: New essays on semantics and pragmatics* (pp. 3-24). Oxford: Oxford University Press.

- Cappelen, H., & Lepore, E. (1997). On an alleged connection between indirect speech and the theory of meaning. *Mind & Language*, 12(3/4), 278-296.
- Cappelen, H., & Lepore, E. (2005). *Insensitive semantics: A defense of semantic minimalism and speech act pluralism*. Malden: Blackwell Publishing.
- Cohen, J. (2013). Indexicality and the puzzle of the answering machine. *The Journal of Philosophy*, 110(1), 5-32.
- Davis, W. (2003). *Meaning, expression, and thought*. Cambridge: Cambridge University Press.
- Davis, W. (2005). *Nondescriptive meaning and reference: An ideational semantics*. New York: Oxford University Press.
- Dever, J. (2013). The revenge of the semantics-pragmatics distinction. *Philosophical Perspectives*, 27, 104-144.
- Ezcurdia, M., & Stainton, R. J. (2013). Introduction. In M. Ezcurdia & R. J. Stainton (Eds.), *The semantics-pragmatics boundary in philosophy* (pp. xiii-xxxiv). Peterborough: Broadview Press.
- Gauker, C. (2008). Zero tolerance for pragmatics. *Synthese*, 165(3), 359-371.
- Grice, P. (1989). *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Horwich, P. (1998). *Meaning*. New York: Oxford University Press.
- Horwich, P. (2005). *Reflections on meaning*. New York: Oxford University Press.
- Johnson, M., & Nado, J. (2014). Moderate intuitionism: A metasemantic account. In A. R. Booth & D. P. Rowbottom (Eds.), *Intuitions* (pp. 68-90). Oxford: Oxford University Press.
- Kaplan, D. (1989a). Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481-564). New York: Oxford University Press.
- Kaplan, D. (1989b). Afterthoughts. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 565-614). New York: Oxford University Press.

- King, J. C. (2013). Supplementives, the coordination account, and conflicting intentions. *Philosophical Perspectives*, 27, 288-311.
- King, J. C., & Stanley, J. (2005). Semantics, pragmatics, and the role of semantic content. In Z. G. Szabó (Ed.), *Semantics versus pragmatics* (pp. 111-164). Oxford: Oxford University Press.
- Korta, K., & Perry, J. (2008). The pragmatic circle. *Synthese*, 165(3), 347-357.
- Korta, K., & Perry, J. (2011). *Critical pragmatics: An inquiry into reference and communication*. New York: Cambridge University Press.
- Kripke, S. (1977). Speaker's reference and semantic reference. *Midwest Studies in Philosophy*, 2(1), 255-276.
- Kripke, S. (2013). *Reference and existence: The John Locke lectures*. New York: Oxford University Press.
- Lepore, E., & Stone, M. (2015). *Imagination and convention: Distinguishing grammar & inference in language*. New York: Oxford University Press.
- Levinson, S. C. (2000). *Presumptive meanings: The theory of generalized conversational implicature*. Cambridge: MIT Press.
- Lewis, D. (1975). Languages and language. In K. Gunderson (Ed.), *Language, mind and knowledge: Minnesota studies in the philosophy of science*, Vol 7 (pp. 3-35). Minneapolis: University of Minnesota Press.
- Loar, B. (1976). Two theories of meaning. In G. Evans & J. McDowell (Eds.), *Truth and meaning: Essays in semantics* (pp. 138-161). London: Oxford University Press.
- Millikan, R. G. (2005). *Language: A biological model*. Oxford: Oxford University Press.
- Montague, R. (1974). Pragmatics. In R. Thomason (Ed.), *Formal philosophy: Selected papers of Richard Montague* (pp. 95-118). New Haven: Yale University Press.
- Perry, J. (2001/2012). *Reference and reflexivity*. Stanford: CSLI Publications.
- Perry, J. (2009). Directing intentions. In J. Almog & P. Leonardi (Eds.), *The philosophy of David Kaplan*

- (pp. 187-201). New York: Oxford University Press.
- Perry, J. (2017). Indexicals and undexicals. In *Reference and representation in thought and language* (pp. 46-56). New York: Oxford University Press.
- Reimer, M. (1991). Do demonstrations have semantic significance? *Analysis*, 51(4), 177-183.
- Salmon, N. (2005). Two conceptions of semantics. In Z. G. Szabó (Ed.), *Semantics versus pragmatics* (pp. 317-328). Oxford: Oxford University Press.
- Sidelle, A. (1991). The answering machine paradox. *Canadian Journal of Philosophy*, 21(4), 525-539.
- Skyrms, B. (2010). *Signals: Evolution, learning, and information*. Oxford: Oxford University Press.
- Speaks, J. (2016). The role of speaker and hearer in the character of demonstratives. *Mind*, 125(498), 301-339.
- Stojnic, U., Stone, M., & Lepore, E. (2017). Discourse and logical form: Pronouns, attention and coherence. *Linguistics and Philosophy*, 40(5), 519-547.
- Szabó, Z. G. (2005). Introduction. In Z. G. Szabó (Ed.), *Semantics versus pragmatics* (pp. 1-14). Oxford: Oxford University Press.
- Wettstein, H. K. (1984). How to bridge the gap between meaning and reference. *Synthese*, 58(1), 63-84.