

This is the pre-peer reviewed version of the following article: Åkerman, J. 'Indexicals and Reference-Shifting: Towards a Pragmatic Approach', Philosophy and Phenomenological Research (forthcoming), which has been published (Early View) in final form at <http://dx.doi.org/10.1111/phpr.12216>. This article may be used for non-commercial purposes in accordance with [Wiley Terms and Conditions for Self-Archiving](#).

Indexicals and Reference-Shifting: Towards a Pragmatic Approach

JONAS ÅKERMAN

Abstract

I propose a pragmatic approach to a range of cases of reference-shifting in indexicals, involving e.g. written notes, answering machines and talk about game pieces. I proceed in two steps. First, I prepare the ground by arguing that the arguments against such a pragmatic approach raised in the recent literature do not provide us with any principled reason for ruling it out. Second, I take a first few steps towards implementing this approach, by sketching a pragmatic theory of reference-shifting in indexicals, and showing how it can handle some of the problematic cases. On the one hand, the immediate scope of the paper is restricted to indexicals and reference-shifting, and the discussion is confined to a specific range of theories and cases. But on the other hand, the general approach proposed is compatible with a fairly broad range of more or less semantically conservative theories, and many of the conclusions drawn are significant for the evaluation of pragmatic explanations in philosophy more generally. In other words, the ideas and arguments presented in this paper have implications that go beyond its most immediate concerns. The overall goal is to offer a new perspective on the issues under discussion, and to prompt philosophers to reconsider some of the established methods by which pragmatic explanations are evaluated.

1. Introduction

In the recent debate on indexical reference, a certain class of examples has had a particularly prominent role to play, namely those involving recorded utterances of various kinds, like written notes and answering machine messages. My main aim in this paper is to bring a novel perspective to this debate, by proposing a pragmatic approach to the kind of reference-shifting exemplified by these cases. Accordingly, the discussion will be confined to a specific range of cases, including those already mentioned. More generally, I will only be concerned with purely indexical uses of indexicals, i.e. uses in which the

Address for correspondence: Department of Philosophy, Stockholm University, SE-106 91 Stockholm, Sweden.
Email: jonas.akerman@philosophy.su.se

referent is dependent on the context of use and does not depend on some associated demonstration (cf. Kaplan (1989, 489–491). This means that, for present purposes, certain cases of reference-shifting are left out, including, of course, those involving demonstrative uses of indexicals,¹ but also those in which the shifts are due to the indexicals being used in free indirect discourse.² But although the scope of the paper is restricted in this way, and although the immediate significance of the cases in question is limited to the recent debate on indexical reference, several of the conclusions to be drawn in the course of the discussion have important implications for other philosophical debates as well, since they concern the evaluation of pragmatic explanations in general. Indeed, a subordinate, but yet important goal of the paper is to prompt philosophers to reconsider some of the established methods for evaluating pragmatic explanations.

The structure of the paper is as follows. In section 2, I start by giving a brief introduction to the standard basic framework for handling indexicals, together with a simple view on how to spell out some of the most essential details. I then present some examples of reference-shifting, which seem to cause trouble for the simple view, and consider some ways in which it might be modified in order to deal with them. Then I present the pragmatic approach to be explored in the rest of the paper. In section 3, I consider objections that have been raised against the possibility of giving a pragmatic account of the problem cases, and conclude that they are insufficient to rule out the pragmatic approach. Having thus prepared the ground, I turn, in section 4, to the task of working out the details of the pragmatic approach, by taking a first few steps towards developing a pragmatic account of the relevant cases of reference-shifting. I conclude with some remarks on the merits of the proposed account.

¹ A standard example of a demonstrative use of an indexical is when someone points at a city on a map and says ‘In two weeks, I will be here’ (Kaplan (1989, 491) cites this example and attributes it to Michael Bennett).

² There are several suggestions in the offing regarding how the phenomenon of free indirect discourse should be analysed. For instance, it may be taken to involve parataxis, metalinguistic demonstratives and context-shifts (cf. Banfield 1973 and Schlenker 2004), or it may be taken to be a form of mixed quotation (Maier 2014). For present purposes, we need not enter into the details of the behaviour of indexicals in indirect free discourse.

2. Reference-shifting and conservative semantics

2.1 Indexicality: a simple view

Natural languages, like English, contain so-called indexical expressions, like ‘I’, ‘now’, and ‘here’, whose contents vary with certain features of the situation in which they are uttered. The referent of ‘I’ appears to depend on the identity of the utterer, the referents of ‘now’ and ‘here’ appear to depend on when and where the utterance takes place.³ So, to take a very simple example, if Amanda were to utter the sentence

- (1) I am so happy to be here now.

in her grandparents’ apartment at 3 p.m. on Christmas Eve 2014, the referent of ‘I’ would be Amanda, the referent of ‘here’ would be Amanda’s grandparents’ apartment, and the referent of ‘now’ would be 3 p.m. on Christmas Eve 2014.⁴ If someone else would have uttered the same sentence, or if Amanda would have uttered it at some other time or place, the referents of these expressions would not have been the same.

The idea that natural languages exhibit this form of context sensitivity is universally accepted among contemporary theorists, and no account of natural language would be regarded as adequate unless it could accommodate this phenomenon. One of the standard ways to do this is to adopt the framework developed by David Kaplan in his seminal paper ‘Demonstratives’ (1989). Kaplan distinguishes meanings of two kinds: *character* and *content*. The character of an expression is invariant across utterance situations, and is taken to represent its purely linguistic or conventional meaning; it is a standing “meaning rule” that allows competent speakers to ascertain the content of the expression as used on a given occasion. The content of an expression as uttered on a given occasion is what the speaker expressed or said by so using it, e.g. the referent of ‘I’ as used by Amanda, or the proposition expressed by her utterance as described above.

The characters of ‘I’, ‘now’ and ‘here’ may thus be thought of as functions that take certain designated features of the utterance situation as arguments, and yield referents as values. Following Kaplan in representing the designated features of the utterance situation as parameters of a context-index, we may represent characters as functions from context-indices to contents as follows

³ As noted above, I will only be concerned with genuinely indexical uses of these expressions.

⁴ Of course, such facts about the production of the utterance would not in general be sufficient to determine precise spatio-temporal boundaries for the referents of ‘here’ and ‘now’, but for present purposes we may ignore this complication.

(where $[[e]]$ is the character of e and C is a variable for context-indices): $[[I]]_C$ =the agent (typically the speaker) of C (a_C); $[[now]]_C$ =the time of C (t_C); $[[here]]_C$ =the location of C (l_C). The context-index C_1 pertaining to Amanda's utterance of (1) would thus look as follows: $\langle a_{C_1}$ =Amanda, t_{C_1} =3 p.m. on Christmas Eve 2014, l_{C_1} =Amanda's grandparents' apartment, ... \rangle .⁵

It is important to remember that context-indices—for which Kaplan uses the term 'contexts'—are distinct from any real utterance situations that might be taken to represent. This allows us to abstract away from actual utterance situations in the formal semantics, which is welcome given Kaplan's main goals. But of course, if we want to apply the framework to real utterances, we need to say something about the relation that need to hold between an utterance and a context-index in order for the latter to count as representing the former. A simple way to specify this relation with respect to the parameters considered above would be to take a context-index C to pertain to an utterance U only if a_C =the producer of U , t_C =the time at which U is produced, and l_C =the place at which U is produced.

2.2 Some problem cases

At a first glance, this simple view may seem to adequately accommodate the phenomenon of indexical reference shifts, and it also seems to give the intuitively correct result in many cases, including the one considered above. Since Amanda produced her utterance of (1) at 3 p.m. on Christmas Eve in her grandparents' apartment, the specification of the representing relation just given entails that the context-index pertaining to her utterance is C_1 , as desired.

On a closer look, however, there seem to be reason to question the adequacy of this simple view. For instance, it looks as if the referent of 'now' can vary between two utterances, even if they are produced simultaneously. Suppose that in answering a request to visit Mrs Horne, her secretary Mark utters the following sentence:

(2) Mrs Horne is busy now.

Moreover, suppose that Mrs Horne is busy recording a message for her answering machine, and thus utters the following sentence, exactly at the same time as Mark utters (2):

(3) I cannot take your call now.

⁵ The dots indicate that this specification of C_1 is not complete. In general, context-indices also contain a parameter for the world of the utterance, and further parameters would be needed in order to handle other indexical expressions, but for present purposes, we do not need to list all of them.

A few hours later, when Mrs Horne has left the office, someone calls her office phone and the answering machine message, including the recorded utterance of (3), is played back. Mrs Horne's and Mark's utterances are identical as regards their time of production, but presumably, the referents that they express by their respective utterances of 'now' are not the same.

A similar case can be constructed for 'here'. Suppose that after leaving the office, Mrs Horne goes to her summer house, where she plans to stay until the next day. The following morning, she sits down at her desk and writes a postcard for her mother, including the following sentence:

(4) It is sunny and warm here today.

She then decides to take advantage of the pleasant weather, and stay for an extra day. Still at her desk, she writes a note saying

(5) I am not here today.

and sends it to Mark together with an instruction to post it on her office door. Mrs Horne's utterances of (4) and (5) are identical as regards their place of production, but presumably, the referent of 'here' shifts between them.

It even seems as if the referent of 'I' can shift across utterances that do not differ with respect to who is producing them. Suppose that Mrs Horne is playing Monopoly with some of her friends, and utters the following sentence in order to inform them about the location of her game piece:

(6) I am on Baltic Avenue.

A few days later, she is meeting a friend in Atlantic City. She calls her up on her mobile phone and utters (6) in order to inform her friend of her present location. The producer of both of these utterances is Mrs Horne, but nevertheless, the referent of 'I' appears to shift across them.⁶

⁶ One may, following Nunberg (1993), argue that there is no genuine reference-shift with respect to 'I' in the Monopoly case, and instead take 'is on Baltic avenue' to express the property of having one's game piece placed on the square marked 'Baltic avenue'. However, Nunberg's arguments have been questioned by Mount (2008), who also argues against a pragmatic approach of the general kind I propose in what follows. I will not try to settle this issue here, but rather, for present purposes, simply assume that this is a genuine case of reference-shifting with respect to 'I'. This assumption will, if anything, make it harder to fulfil my aim of showing that one can handle these apparent counterexamples while sticking to a conservative (meta)semantics. (An observation that seems to support this assumption is that it seems acceptable to add 'and so is that shoe-shaped game piece' to (6) in the Monopoly case,

In the light of these examples, there is some reason to suspect that the simple view sketched initially is deficient. It seems unable to account for the shifts in reference that seem to occur in the cases described, and it yields the intuitively wrong result in half of them.

As regards the cases involving ‘now’ and ‘here’, the culprit appears to be the idea that only facts about the production of the utterance are semantically relevant. A first step towards remedying the simple view would be to acknowledge the semantic relevance of facts about the *tokening* of the utterance. This is precisely what Jonathan Cohen (2013) proposes that we do, and, primarily motivated by cases involving answering machines, he suggests (in effect) that we should take t_c to be the time at which U is tokened, and l_c to be the place at which U is tokened.

Cohen’s tokening view is intended as a conservative extension of the simple view; it only departs from it in cases when the production and tokening of the utterance are temporally or spatially distinct. This means that the tokening view retains many predictions of the simple view, e.g. the one regarding Mark’s utterance of (2). In contrast, when it comes to Mrs Horne’s utterance of (3), the time of production and the time of tokening come apart. As a result, the tokening view predicts, correctly it seems, that the referent of ‘now’ is the time of playback, while the simple view yields the counterintuitive prediction that it is the time of recording. The same point holds for Mrs Horne’s utterance of (5): the place of production and the place of tokening comes apart, and the correct prediction seems to be that ‘here’ as it occurs in this utterance refers to the place of tokening, i.e. Mrs Horne’s office, rather than the place of production, i.e. Mrs Horne’s summer house.

However, there are also cases in which the simple view seems to get it right, even though production and tokening come apart in the relevant respect. For instance, when it comes to Mrs Horne’s utterance of (4), the intuitively correct prediction is that ‘here’ refers to the place of production, i.e. Mrs Horne’s summer house, rather than the place of tokening, i.e. the place at which her mother reads the postcard. Assuming that the postcard is not read the same day as it was written, the same point holds for ‘today’ as it occurs in Mrs Horne’s utterance of (4): intuitively, the referent is the day on which the letter was written, not the day on which it is read. So, while adopting the tokening view may be a step in the right direction, it does not seem flexible enough to accommodate the observation that utterances of ‘here’ and ‘today’

even though it makes no sense to ascribe the property of having one’s game piece placed on the square marked ‘Baltic avenue’ to a game piece. For more examples, and discussion, see Mount (2008).)

can sometimes be used to refer to the place/day of production even when it does not coincide with the place/day of tokening.⁷

One way to get a more flexible account would be to modify the semantic (or metasemantic⁸) part of the theory. For instance, one could follow a recent suggestion by Eliot Michaelson (2014), according to which the features associated with t_c and l_c vary as a function of the type of situation in which the utterance is made. In face-to-face type situations, t_c =time of production and l_c =location of production; in answering-machine type situations t_c =time of playback and l_c =location of playback; in postcard type situations t_c =time of production and l_c =location of production. This yields an account that predicts the intuitively correct results for the cases involving ‘now’ and ‘here’.

Of course, the price for making this move is a more complicated semantics, and moreover, it does not seem to be quite enough in order to handle all the problem cases. In particular, it does not seem to help at all with the example involving the two utterances of (6). Neither the production-tokening distinction, nor the situation-type relativity appealed to above seems sufficient to explain the shift in referent between these cases. Indeed, shifts like this can occur even within one and the same utterance. Consider the following sentence as uttered by Mrs Horne during a game of Monopoly:

(7) I know that I am on Baltic Avenue.

Both occurrences of ‘I’ are uttered in the very same face-to-face situation, and the tokening and production of the utterance coincide in all possible respects. Yet, the referent of ‘I’ appears to shift from Mrs Horne in the first occurrence to her game piece in the second.

2.3 The pragmatic approach

So, even if it seems possible to make significant progress when it comes to accommodating initially problematic cases by tinkering with the semantic (or metasemantic) theory, the moves considered so far do not seem to take us all the way. At this point, it one may be tempted to take a further step in complicating the relation between utterance situations and context-indices, by taking it to depend on the speaker’s intentions. This would allow us to take Mrs Horne’s intentions into account in determining the semantic value of the expressions she used, and since she can plausibly be assumed to have intended the referent of the first occurrence of ‘I’ to be herself, and the referent of the

⁷ For further discussion of this case, see Cohen (2013) and Michaelson (2014).

⁸ As I am using the term ‘metasemantics’ here, the metasemantic part of the theory concerns the relation between formal representations and what they represent, e.g. the relation between real-life utterance situations and context-indices.

second to be her game piece, we can thus get the intuitively right result for this case as well.

While a number of authors have defended different versions of such an intentionalist view in recent years, it remains highly controversial, and faces several serious objections.⁹ I cannot go into details about this here, but suffice it to say that it remains unclear whether these objections can be met. In addition, there is some reason to resist intentionalism on more general theoretical grounds. Keeping the semantics (and metasemantics) free from intentions is one way of satisfying the simplicity desideratum, and thus it may be argued that the complications that come with intentionalism are problematic in their own right, and should be avoided if possible.¹⁰ In the rest of this paper, I will explore this possibility. More specifically, I will explore the prospects of giving purely pragmatic explanations of the kinds of reference-shifting exemplified in the cases considered above.

One interesting question concerns the scope of this strategy. Could it be used to account for all of these cases, and if so, would this suffice to save even the simple view from these apparent counterexamples? Perhaps the complications introduced by relatively conservative extensions of it, like the ones suggested by Cohen, could be avoided as well? Interestingly, Cohen himself emphasizes the importance of distinguishing between “semantically relevant” intuitions, and intuitions that track merely pragmatically conveyed content. According to Cohen, we should be careful neither to reject nor accept theories on the grounds of intuitions that are not of the appropriate kind:

After all, the theory that the answering-machine cases threaten is a theory about semantic reference. As such, cases shouldn't be counted as counterexamples to a candidate semantic theory unless we have reason to think that, in those cases, specifically semantic reference behaves in a way that diverges from the predictions of the theory. (Cohen, 2013, 7–8)

However, Cohen argues that we do have reasons to take the intuitions about answering machine cases to track semantic content, and expressed scepticism about the prospects for a pragmatic treatment. This attitude is not uncommon.

⁹ For defences of intentionalist views of different kinds see Predelli (1998a, 1998b), Åkerman (2009, 2010), Stokke (2010), Montminy (2010), King (2014). For criticism see Gorvett (2005), Romdenh-Romluc (2006), Gauker (2008), Cohen (2013), Cohen & Michaelson (2013).

¹⁰ For instance, given the multitude and complexity of the factors that determine speakers' intentions, adopting intentionalism would seem to commit us to what Glanzberg (2007, 25) calls “a highly indirect metasemantics”, which in turn would give us a messier overall theory.

On the contrary, the pragmatic approach has, insofar as it has been seriously considered at all, met with disapproval, and several objections has been raised against the possibility of giving a pragmatic account of cases like the ones considered above. In order to show that the pragmatic approach is worthy of further consideration, it must be shown that these objections can be met. Thus, a significant part of this paper will be devoted to this task.

3. Arguments against the pragmatic approach

3.1 The answering machine puzzle

Like many authors, Cohen focuses on a case that seems particularly problematic for the simple view, since it gives rise to the so-called *answering machine puzzle*.¹¹ Suppose that Mrs Horne had recorded the following sentence as part of her answering machine message:

(8) I am not here now.

The simple view does not only get the reference of ‘now’ intuitively wrong for this message, but it also predicts that it expresses that the producer of the utterance was not at the place of production at the time of production. This would be an absurd thing to say, insofar as it lies in the nature of utterance production that the producer of an utterance must be at the place of production at the time of production. Moreover, this is not what we would normally take an answering machine occurrence of (8) to say.

At this point, it would be natural for a proponent of the simple view to appeal to a pragmatic explanation, and claim that our intuitions about this case do not in fact track semantic content, but rather what is merely pragmatically conveyed. However, Cohen argues that this strategy is not viable. His argument is based on two familiar tests for distinguishing between semantic and pragmatic content, both of which are due to H.P Grice:¹²

Suppose I record/inscribe an answering-machine message containing ([8]) in context c_1 , you call later when I am at home, but I screen the call and allow the outgoing message to play its token of ([8]) in c_2 . Crucially, the (false) content expressed in c_2 does not seem detachable or cancelable in

¹¹ This puzzle was introduced by Sidelle (1991), under the name ‘the answering machine paradox’.

¹² See the papers collected in Grice (1989), in particular ‘Logic and conversation’ and ‘Further notes on logic and conversation’. These tests will be considered in more detail below.

the way that we expect of extrasemantic content. The best account, then, is that ([8]) semantically expresses something false in c_2 iff the agent is located at the time/place of c_2 (hence true in c_2 iff the agent is not located at the time/place of c_2). (Cohen, 2013, 8n)

The argument here could be restated as follows: If the conveyed content in question is not semantically expressed, it must be either conventionally or conversationally implicated. If the former, it should be detachable, if the latter, it should be cancellable.¹³ Since it is neither, it is neither conventionally nor conversationally implicated, and thus it must be semantically expressed, rather than merely pragmatically conveyed.¹⁴

Michaelson (2014) provides a similar, but somewhat more detailed argument, based on the cancellability test. He targets the suggestion that answering machine recordings of (8) conversationally implicate a content that is distinct from the one assigned by Kaplan's theory, or, more specifically, that "while answering machine tokens of ([8]) semantically express that the speaker isn't present at the moment of utterance, they pragmatically convey (via a conversational implicature) that she isn't present at the time and place of playback" (2014, 535). The argument is neat and simple: the putative conversational implicatures do not seem to be cancellable. If they were, Michaelson argues, the following sentence as played back on an answering machine would be non-contradictory "since the first and second clause would be about different times, and possibly about different places as well" (2014, 536):

(9) I am not here right now, but Eliot is at home at the time when you're calling.

However, since (9) seems internally contradictory, "recorded tokens of ([8]) do not merely conversationally implicate that they are about the place and the time of playback" (2014, 536).

¹³ I will say more about these kinds of implicatures shortly.

¹⁴ An obvious problem with the argument so understood is that it assumes that if a content is merely pragmatically conveyed, it is either conventionally or conversationally implicated. While this is a questionable assumption, my reply to Cohen's argument does not depend on rejecting it. However, one could also understand Cohen's argument as being independent of this assumption. Instead, one could take it to rely on the assumption that any merely pragmatically conveyed content must either be detachable or cancellable. This assumption may also be questioned, but for the purposes of my response, I can happily go along with this assumption as well.

I think that both of these objections fail, even if we go along with the assumption that the contents in question must pass either the detachability test or the cancellability test in order to count as merely pragmatically conveyed. In fact, I think that *pace* Cohen and Michaelson, they are cancellable, albeit not always felicitously so (I will elaborate this shortly). As regards the detachability test, I think it is a mistake to expect that it could do any work in this dialectical context.

3.2 The argument from detachability

According to the standard Gricean notion of detachability, an implicature is detachable if there is some way of saying the same thing that does not carry the implicature (Grice, 1989, 39). More generally, a content C conveyed by an utterance U of a sentence S is detachable if there is some other sentence S^* that can be used to say what U says without thereby also conveying C . For instance, the implicature generated by (10) is detachable, since what is said by (10) can be said by (11) without generating that implicature:

- (10) She is poor but honest.
(11) She is poor and honest.

The content conveyed by (10), but not by (11), is a standard example of a so-called conventional implicature, i.e. an implicature that turns on the conventional meaning of the words used. In general, such implicatures are characterized by their being detachable (in contrast to semantic entailments) but not cancelable (in contrast to conversational implicatures; more on this below).

Turning to Cohen's example, the question is whether or not the putative pragmatic content conveyed by (8) in c_2 is detachable in the same way as the implicature generated by (10). Cohen claims that it is not, but as this issue crucially turns on what (8) as played back in c_2 says or semantically expresses, it seems hard to establish this claim on sufficiently independent grounds. Let me elaborate.

According to the simple view, (8) as played back in c_2 semantically expresses the following:

- (12) *Dthat*[the speaker] is not at *Dthat*[the place of the recording machine] at *Dthat*[the time of recording].¹⁵

¹⁵ The operator '*Dthat*' was introduced by Kaplan (1978; 1989, 521–522), and defined as “the demonstrative ‘that’ with the following singular term functioning as its demonstration”. In other words, for any singular term t , *Dthat*[t] is a directly referential term referring to the

If this is right, (12*)—in which PR and TR are names for the place and time of recording, respectively—could be used to say what (8) says in c_2 :

(12*) I am not at PR at TR.

Such an utterance of (12*) would not convey that the speaker is not at the place of the recording machine at the time of playback. So, according to the simple view, this content, insofar as it is conveyed by (8) as played back in c_2 , is detachable.

According to Cohen's theory, (8) as played back in c_2 semantically expresses the following:

(13) *Dthat*[the speaker] is not at *Dthat*[the place of the recording machine] at *Dthat*[the time of playback].

If this is right, there is clearly no way of saying what (8) as played back in c_2 says without also expressing (13). So according to Cohen's view, this content is not detachable.

What this shows is that the result of the detachability test turns on what (8) as played back in c_2 semantically expresses. Since this is precisely what the disagreement between Cohen and a proponent of the simple view consists in, neither of them is in a position to appeal to the test in any dialectically efficient way.¹⁶ In particular, the possibility of sticking to the simple view while giving a pragmatic account of cases involving recorded messages cannot be excluded by appeal to this test. More generally, we should not expect the detachability test to be of much use when it comes to settling the score between competing semantic accounts of indexicals. Thus, that part of Cohen's objection can be put aside. Instead, we now turn to the cancellability test.

denotation of *t*. This device is helpful here as it makes clear that the propositions expressed here are singular rather than general.

¹⁶ In fact, the same point can be made with respect to (10). Although this is standardly used to exemplify conventional implicature, it is not entirely uncontroversial. For instance, Kent Bach (1999) argues that that uttering (10) instead of (11) does not merely generate an implicature, but actually makes a difference to the semantic content expressed. A dispute of this kind cannot be settled by appeal to the detachability test, since that test itself would turn on whether (10) and (11) in fact express the same semantic content. As Bach puts it, "detachability is not an independent test" (Bach, 1999, 330).

3.3 The argument from cancellability

According to Grice, cancellability is a necessary condition for being a (mere) conversational implicature. He distinguishes between two forms of cancellability:

[A] putative conversational implicature that p is explicitly cancelable if, to the form of words the utterance of which putatively implicates that p , it is admissible to add *but not p* , or *I do not mean to imply that p* , and it is contextually cancelable if one can find situations in which the utterance of the form of words would simply not carry the implicature. Now I think that all conversational implicatures are cancelable [...]. (Grice, 1989, 44)

Given the assumption that “extrasemantic”¹⁷ content in general is either cancellable or detachable, a defender of the simple view would have to show that the content conveyed in answering machine occurrences of (8) passes at least one of these tests in order to be in a position to appeal to a pragmatic explanation of our intuitions about this case. Having put the detachability test to one side, it remains to evaluate the situation with respect to the cancellability test.¹⁸

Cohen’s claim is that in a “screening” case—i.e. a case in which the person P who recorded the message in context c_1 is standing beside the answering machine at the time when the message containing (8) is played back in context c_2 —the purported implicature is not cancellable, and this is supposed to show that it cannot really be a (mere) conversational implicature.¹⁹ A problem with

¹⁷ A (perhaps obvious) terminological point: since the phrase ‘extrasemantic content’ as used by Cohen in this context is meant to cover conventionally implicated material, the corresponding use of ‘semantic content’ cannot be taken in a sense that counts conventional implicatures as part of semantic content (cf. Barker 2003, 3).

¹⁸ It should be noted that the above arguments focus on explicit rather than contextual cancellability, and that there seem to be good reasons to do so. What would be needed in order to show that the content conveyed by answering machine occurrences of (8) is contextually cancellable in any interesting sense would be a situation in which all parties could agree that an utterance of (8) would semantically express the same content as in the answering machine case, but still would not convey what it conveys in the answering machine case. The prospects for this look rather dim (due to considerations of the kind raised in connection with the discussion of detachability). For this reason, the present discussion will be restricted to explicit cancellability.

¹⁹ Note that the fact that the message will convey a false content even if the speaker tries to cancel the putative implicature is a red herring, as it cannot be taken to show that the cancellation fails. On the simple view, the content conveyed in c_2 will be false regardless of whether the cancellation of the putative implicature succeeds or not. If it succeeds, the content

this is that even if it appears impossible for the speaker in this case to cancel the purported implicature, this apparent incancellability may well be due to some special feature of this rather peculiar case. For instance, it is far from obvious that speakers should in general be expected to be in a position to cancel implicatures generated by previously recorded utterances. Suppose that someone were to record Mrs Horne when she answers a question about Mark's administrative abilities by uttering the following sentence (and nothing more),

(14) Well, he is quite handsome.

thus conversationally implicating that Mark is not a very good administrator. Moreover, suppose that this conversation were to be played back in the presence of a group of people including Mrs Horne herself. It seems rather dubious that she would be in a position to cancel the implicature in the situation in which the recording is played back.²⁰ Nevertheless, this is an uncontroversial instance of conversational implicature. If this is right, then, why should the apparent fact that the speaker in the screening case is not in a position to cancel the corresponding content count against it being merely conversationally implicated?

Although a discussion about the special features of the screening case alluded to above might shed some further light on the general conditions for cancellability, I will not pursue this issue here. My purpose here is merely to show that the cancellability test does not count against a pragmatic explanation of answering machine occurrences of (8), and, as will become clear shortly, that point can be made independently of such a discussion. Instead, I will turn to the more straightforward argument offered by Michaelson.

As we saw above, Michaelson appeals to examples like the following in order to argue that the putative conversational implicature fails the cancellability test.

(15) I am not here now, but I will be at home at the time when you call.²¹

conveyed will be something like (12), which is false since the speaker (P) obviously was at the place of the recording machine at the time of recording. If it fails, the content conveyed will be something like (13), which is false since the speaker (P) is standing behind the machine, screening the message.

²⁰ At least given that it is part of the common ground in the playback situation that the recorded message indeed contains everything that she said in response to the question.

²¹ I have taken the liberty of substituting 'P' for 'Eliot' in the second clause of Michaelson's example, and modified the tense in the second clause. This makes the sentence sound less odd

To bring out his point, he contrasts (15) with (16), which he takes to exemplify a genuine conversational implicature and a successful cancellation:

- (16) Laura and Danny had a kid and got married, but not in that order.²²

Now, if the content standardly conveyed by (8) (the first clause of (15)) as played back on an answering machine were a genuine conversational implicature, should we not expect it to be cancellable in the same way as the merely implicated content of the first clause of (16)? A natural first-blush reaction would be to answer this question in the positive, and if we do, Michaelson's argument seems hard to resist. Here is the argument in full:

An utterance of just the first clause of [(16)], 'Laura and Danny had a kid and got married', standardly communicates that they did so kid-first. However, there is nothing contradictory about an utterance of [(16)]—which Grice takes to indicate that this ordering is only conversationally implicated by that utterance. Thus, nothing about the order of these events is part of the semantic content expressed by an utterance of [(16)]. On the other hand, when [(15)] is played back on an answering machine, it seems not only to be false, but also internally contradictory. If this is right, then recorded tokens of [(8)] do not merely conversationally implicate that they are about the place and time of playback. If they did, then, by parity of reasoning, playback tokens of [(15)] should be non-contradictory—since the first and second clause will be about different times, and possibly about different places as well. (Michaelson, 535–536)

I agree with Michaelson that (15) differs from (16) in that the former has an air of contradiction that is completely lacking in the latter (as long as we restrict ourselves to answering-machine occurrences). Or rather, I agree that the addition of the second clause in (16) is admissible in a way that the addition of the second clause in (15) is not. Strictly speaking, this means that the (explicit)

without affecting the main point (the pronoun and the proper name in the original example are intended to refer to the same person).

²² It may be worth noting that although (16) involves a conversational implicature on Grice's view, relevance theorists take it to involve an "explicature", which is part of what is said rather than implicated.

cancellability test as Grice formulated it (see quote above) fails for answering machine occurrences of (8) in a way that it does not fail for utterances of the first clause of (16). Nevertheless, I think that the argument can be resisted.

3.4 Infelicitous cancellation

Even on the assumption that the contents standardly conveyed by the first clauses of (15) and (16) are both generated via a conversational implicature, we should not expect answering machine utterances of the former to be felicitous just because standard utterances of the latter are. Rather, I suggest that from the perspective of the simple view, we should view the answering machine utterance of (8) as an instance of the more general phenomenon that it is possible to pragmatically convey something coherent by uttering something that is literally incoherent. The following conversation provides a fairly clear example of how a literally incoherent sentence can be used to pragmatically convey something coherent, and how the utterance as a whole can be rendered infelicitous by way of a successful cancellation:

- (17a) Are you happy with your new job?
(17b) Yes and no.
(17c) But I do not mean to imply that I am happy with it in some respects, and not happy with it in others.²³

The utterance of (17b) implicates that the speaker is happy with the new job in some, but not all respects. This implicature is then cancelled by the addition of (17c), and then the only thing that remains to be conveyed is an outright contradiction.²⁴ In other words, the successful cancellation of the implicature renders the utterance as a whole (17b+17c) infelicitous. In this sense, the implicature is essential to the felicity of (17b).²⁵

Similarly, from the perspective of the simple view, the purported implicature generated by an answering machine utterance of the first clause of (15) is essential to the felicity of the utterance as a whole: if it is cancelled, only the strict semantic content—the absurdity represented in (12)—will remain to be conveyed. In other words, if the simple view is correct, we should expect answering machine utterances of (15) to be infelicitous even if the cancellation

²³ I owe this example to Eric Johannesson.

²⁴ Of course, I do not mean to suggest that an utterance of (17b) in response to an utterance of (17a) will invariably have this, and only this implicature. But we can surely imagine a case in which it would, and that's enough for present purposes.

²⁵ I discuss this phenomenon at length, and provide additional examples, in Åkerman (Forthcoming II).

succeeds. This means that the apparent infelicity of such utterances cannot be taken to show that the purported implicature is not cancellable.²⁶

Of course, this does not show that the content in question is cancellable. It merely shows that we cannot rule this out on the basis of the infelicity of answering machine occurrences of (15), and thus that the standard explicit cancellability test—which turns on felicity intuitions—cannot be applied in the relevant way to answering machine occurrences of (8). What we can do, however, is to apply the test to answering machine utterances involving ‘now’, whose purported implicatures are not essential to their felicity, like Mrs Horne’s utterance of (3):

(3) I cannot take your call now.

When it comes to this case, it does seem that the purported implicature can be cancelled without infelicity. Consider the following sentence as recorded as part of an answering message:

(18) I cannot take your call now [as I am busy recording this message] but I will be able to take your call when this message is played back [since I always screen incoming calls whenever the answering machine is on].

The bracketed material can be either explicit in the message, or part of the common ground in the playback situation. Either way the message comes out as felicitous, and the content *Mrs Horne cannot take your call at the time of playback*, which is purportedly implicated by (3), is clearly not conveyed by (18). In other words, the purported implicature generated by answering machine messages containing (3) appears to pass the explicit cancellability test.

²⁶ According to the analysis offered here, the inadmissibility of adding a cancellation clause to answering machine occurrences of (8) is not to be explained in terms of a contradiction between the clauses. In reply, one could perhaps dig one’s heels in and insist that it is just obvious that the clauses contradict each other. The problem with this reply is that in order for it to be obvious that the clauses contradict each other in the relevant sense, it must be obvious that what the first clause says is contradicted by what the second clause says. Since it is obvious enough that the second clause says that the speaker will be at home at the time of the call (or, that the speaker is at the place of playback at the time of playback), this means that it must be obvious that the first clause says (semantically expresses) that the speaker will not be at home at the time of the call. But to take this as obvious would be question begging. This takes us back to the kind of problem discussed in connection with the detachability test, namely that a dialectically efficient argument cannot, in the context of the present discussion, rely on claims about what answering machine occurrences of (8) semantically expresses (unless they are independently supported).

As cancellability is at most a *necessary* condition for conversational implicature—or perhaps for a somewhat wider class of pragmatic mechanisms—this still does not show that the contents standardly conveyed by answering machine occurrences like (3) or (8) are merely pragmatically conveyed. However, together with the point made about detachability above, this is enough to establish that the standard Gricean tests invoked by Cohen and Michaelson are not sufficient in themselves to rule out the option of sticking with the simple view and explain the problematic intuitions about (3) and (8) in terms of a pragmatic theory.

The general moral to be drawn is that we need to be careful in applying tests like these in cases involving expressions whose semantics are a matter of controversy. While they can give us useful indications about whether a certain content is to be considered as semantically expressed or merely pragmatically conveyed, their applicability is more limited than is usually acknowledged, and unless we take the necessary precautions, they might just as well lead us astray. The significance of this conclusion is not restricted to the present debate. On the contrary, it is pertinent to any philosophical debate in which these tests might be used in order to assess the success of a pragmatic account.

It bears emphasizing that my purpose here is not to defend any particular (meta)semantic theory of indexicals. And of course, the conclusion that the possibility of supplementing the simple view with a pragmatic account of the problem cases can withstand the arguments from detachability and cancellability is not enough in itself to show that it is tenable. Nevertheless, its alleged inability to handle answering machine cases is a particularly oft-cited reason for rejecting it,²⁷ so the possibility of giving an adequate response to this objection is surely relevant to the prospects of its defence. In general, whenever a particular (meta)semantic theory faces counterexamples that might be handled by a supplementary pragmatic theory, the possibility of developing such a pragmatic theory is relevant to its evaluation, at least insofar as there are no other viable strategies available. That the above discussion has focused on the simple view and the answering machine puzzle seems in order given the central place it has occupied in the recent, but in what follows it will become clear that the general pragmatic approach I propose is available to a wider range of theorists. But, again, my purpose is not to defend any of them in particular, but rather to offer a new perspective on the relevant cases and their significance for the evaluation of (meta)semantic theories of indexicals.²⁸

²⁷ See e.g. Sidelle (1991), Predelli (1998a, 1998b, 2011), Cohen (2013), Cohen & Michaelson (2013).

²⁸ For what it is worth, I do not think the simple view is plausible enough in itself to merit acceptance by default. As pointed out above, it seems to lie in the nature of utterance production that the producer of an utterance must be at the place of production at the time of

3.6 Automatic indexicals and linguistic competence

Even if we were to abandon the simple view, simplicity considerations would still seem to motivate a certain degree of conservativeness. In particular, many theorists have, as already pointed out, been very reluctant to bring speaker intentions to bear on the semantic contents of indexicals. This holds in particular for ‘I’, which has been taken as the most obvious example of an “automatic” indexical, i.e. one whose meaning (character) determines a referent purely on the basis of objective facts about the utterance situation.²⁹

As we have seen, a challenge for such conservative views is that there are cases involving ‘I’ that seem difficult for such approaches to handle, and which *prima facie* seem to suggest an intentionalist treatment. A pragmatic approach to these cases would offer an alternative way of handling them, without giving in to intentionalism. Such an approach would also allow us to stick to the idea that ‘I’ is an automatic indexical. However, just like the pragmatic approach to

production. This, in turn, seems to commit the simple view to a restriction on context-indices, according to which sentences should only be evaluated relative to *proper* context-indices—i.e. context-indices C such that a_C is at l_C at t_C (at the world of C). Kaplan endorsed the propriety restriction, and an important motivation for him was that he wanted his theory to capture certain logical relations between indexical expressions, e.g. the relations in virtue of which utterances of ‘I am here now’ and ‘I exist’ appear to be invariantly true. Thus, for Kaplan, the fact that these sentences do indeed come out as true relative to any proper context-index counted in favour of adopting the propriety restriction. There are also passages in Kaplan (1989) that suggest that he had more general reasons for excluding improper context-indices. Improper contexts are, according to Kaplan, “like impossible worlds” in that they could not represent situations that could possibly be realized, and thus they are not interesting for the purposes of semantics (Kaplan, 1989, 509). A more recent endorser of the propriety restriction is François Recanati (2010, 184). Although Recanati applies the restriction to a notion of context which is quite different from that of Kaplan’s notion of a context-index, his motivation is in some ways similar to Kaplan’s. Recanati takes the restriction to be a natural consequence of his idea that a context is “a concrete situation with a particular individual in it endowed with complex mental states” (2010, 184). However, once we look beyond ordinary face-to-face communication, and consider examples like the ones above, the logical relations between indexical expressions that motivated Kaplan’s endorsement of propriety do not seem to hold. Moreover, the propriety restriction does not seem essential to either Kaplan’s or Recanati’s overall view. In Kaplan’s case, it is introduced as a stipulation rather than as a consequence of some core feature of his theory. Indeed, in a footnote, he even acknowledges that the referent of ‘now’ as occurring in answering machine messages might be taken to be the time of playback rather than the time of recording (Kaplan, 1989, 491n). In Recanati’s case, it just seems as if his commitment to the propriety restriction is a result of his having temporarily overlooked the possibility of including tokenings as well as productions of utterances in his contexts, a possibility that he appeals to in other writings (e.g. Recanati 1995), and which does not seem at all incompatible with his “realistic” notion of context (something that Recanati has confirmed in p.c.).

²⁹ The distinction between “pure” and “impure” indexicals was originally drawn by Kaplan (1989, 489–491).

answering machine messages discussed in the previous section, this approach to reference-shifting of the kind exemplified by the different utterances of (6) considered above has met with more criticism than endorsement, insofar as it has been considered at all. In what follows, I focus on the arguments offered by Allyson Mount (2008), and relate them to a general worry raised by Michaelson, which I did not address in the preceding discussion. I will argue that all of these objections can be met, and thus that neither of them provide any principled reason for why a pragmatic explanation of reference-shifts in indexicals should not be included among the options to be seriously considered.

Similarly to Cohen, Mount acknowledges the possibility of taking the intuitions in cases of reference shift for ‘I’ to track merely pragmatically conveyed content, and of taking the competence involved in interpreting the different utterances of (6) to go beyond specifically linguistic or semantic competence. In particular, she notes that if the apparent reference shifts can be explained in terms of conversational implicature, these cases do not really count as genuine counterexamples to the thesis that ‘I’ is an automatic indexical. It should be noted, though, that her arguments against a pragmatic treatment of these cases are more general than these remarks suggest, as they do not turn on any particular features of conversational implicature.

Mount’s strategy is to try to show that understanding the relevant reference-shifting uses of ‘I’, like Mrs Horne’s utterance of (6) while engaged in a game of Monopoly, does not require being “fully conversationally-competent”, but merely being competent in a “specifically linguistic sense” (2008, 205). She offers two general reasons for thinking this:

- Reference-shifting uses of language (in general) are standard in some “official” contexts, like board game instructions.
- Appropriate answers to open-ended questions often require assuming that ‘I’ is not an automatic indexical. (Mount, 2008, 205)

I will consider these in turn.

3.7 The argument from “official contexts”

The idea behind Mount’s first point is that since reference-shifts of the relevant kind occur in rules for board games (like *Monopoly* or *Wise and Otherwise*), there is reason to think that “understanding reference-shifting isn’t just some extra ability that most speakers have; it’s necessary in order to understand the simple instructions of a children’s game” (2008, 206). Since these rules “are presumably designed to be clear, unambiguous, and literal”

(Mount, 2008, 207), our understanding of them should not be taken to rely on anything but specifically linguistic competence. Hence, a pragmatic explanation of these reference-shifts is ruled out.

I do not find this line of argument very convincing, for the simple reason that it does not seem at all unusual for people to rely on more than prospective interpreters' specifically linguistic competence in designing instructions intended for a broad, normally competent audience.³⁰ Consider the following sentence as inscribed on a sign in an elevator:

- (19) If the elevator stops between two floors, press the emergency button.³¹

This is a real-life example of an instruction that is obviously intended to be as clear and simple as possible. Still, understanding it arguably requires more than merely linguistic competence, since a literal understanding would amount to taking it as an instruction to press the emergency button whenever the elevator stops between two floors—for instance when it stops at floor 2, located between floor 1 and 3. Someone who interpreted the message in this way would arguably have misunderstood it at a fairly basic level. Yet, we are not thereby forced to explain such a misunderstanding in terms of a failure to decode the message semantically. On the contrary, the subject in this case appears to have done precisely what the message would be taken to say on an obviously unintended but perfectly literal interpretation. The misunderstanding is most plausibly traced back to some defect in the subject's pragmatic abilities and/or her background knowledge (the two of which are obviously related).³²

This example shows that grasping pragmatic content is sometimes very basic to understanding an utterance. It also shows how quick and automatic this grasp can be. Most ordinary competent speakers would not even notice that their default interpretation in cases of this kind goes beyond what the sign literally says. These are hardly original or novel points, but they are highly

³⁰ It seems that even legislators sometimes rely on interpreters' invoking more than specifically linguistic competence (perhaps mistakenly). For instance, see Stephen Neale's (2007, 251–252) discussion of *Smith v. United States* (1993).

³¹ I owe this example to Peter Pagin.

³² The relevant part of the conveyed content also seems clearly cancellable, but this is not my reason for taking it to be pragmatically rather than semantically conveyed. Indeed, this reason would be insufficient, since, as pointed out by e.g. Saddock (1978, 293) and Davis (1998, 46), cancellability does not distinguish between conversational implicature and ambiguity. Perhaps a case could be made for the claim that 'between' is ambiguous in the relevant way; I shall not address this issue here. Another possibility would be to treat this as a case of relevance-theoretic enrichment, but then it would still involve a pragmatic process (even if truth-conditionally intrusive).

pertinent to the issues discussed here in that they serve to show that there is no general reason to think that pragmatic accounts cannot explain phenomena which occur regularly and frequently, and concern basic aspects of interpretation.

These observations also serve to address a general worry about pragmatic explanations raised by Michaelson in connection with answering machine cases, namely that “pragmatic accounts look bound to offer us poor explanations of the sorts of strong default interpretations—norms of interpretation even—that we see associated with tokens of ‘here’ and ‘now’ as recorded on answering machines” (2014, 536). Michaelson may well be right to hold that failure to interpret an answering machine token of (8) in the standard way would be considered to show that one has “*failed to understand* the sentence-token at a fairly basic level” rather than just failed “to understand all the exigencies of how that utterance or recording token was being used in some particular context” (536). Familiarity with certain types of context—e.g. standard answering-machine contexts—may well be something that we usually expect from competent speakers, and such understanding may be required in order to live up to basic “norms of interpretation”. However, the interpretative abilities that we normally expect competent speakers to have seem in general to go beyond what may plausibly be taken as falling under purely linguistic or semantic competence.

3.8 The argument from open questions

Let us turn to Mount’s second objection to the pragmatic approach, i.e. that appropriate answers to open-ended questions often require assuming that ‘I’ is not an automatic indexical. Mount claims that making sense of certain conversations requires assuming that ‘I’ is not an automatic indexical:

For instance, someone could ask “Who are you?” in a situation (like Monopoly-playing) where “I’m the shoe” would be an appropriate answer and “I’m Allyson” would be inappropriate. Similarly, someone could ask “Where are you?” when all the humans in the room were clearly visible but their gamepieces (or cars, or whatever) were not. By using the second person pronoun in an inquiry, the questioner sets the respondent up to reply using the first person to pick out something other than the speaker. There is no question about the respondent having odd or idiosyncratic intentions in doing so, since she isn’t the one initiating the discourse. And in these situations, using ‘I’ to pick out an object is not just one type of appropriate response, it’s the *only* type of

(non-sarcastic) appropriate response. It's the type of response one must understand in order to make sense of the discourse. (Mount, 2008, 206)

I do not think that this is very convincing either. While Mount is certainly right as regards the appropriateness of the different uses of 'I' in the situations described, this does not show much about the semantic content of the pronoun as so used. As she herself acknowledges, it is not uncommon to say things that are literally false in order to convey things that are true, and sometimes the only appropriate way to (non-sarcastically) convey a certain content may be to say something that is literally false. Mount's reason for thinking that this is not a case in point is that it would be "strange to think that [the speakers in these situations were] forced into saying something false in order to answer the question appropriately" (2008, 207). But why would this be strange, and in what sense are they "forced"? According to Mount, they could not give the purportedly true answers without misunderstanding what they were asked or "[violating] conversational norms and [coming] off as rude or sarcastic" (207). Again, this seems right, but why think that violating conversational norms or coming off as rude or sarcastic is in any way incompatible with saying what is literally true? One thing that can be particularly annoying with people that do behave in this way is when they, after having been reproached for being rude, difficult, or whatever, reply by saying something like 'But what do you mean? I *am* Allyson!' thus trying to justify their behaviour by (correctly, it seems) insisting that they were just telling the truth. The natural way to express one's disapproval of such behaviour is not to say something like 'That is false', but rather to say something like 'Stop being such a literal-minded pain in the neck. You know what I meant'.

Let us take stock. I have argued that the general reasons offered in the recent literature against the possibility of giving a pragmatic explanation of different cases of reference-shifts involving indexicals like 'I', 'now' and 'here', are not convincing enough to motivate an exclusion of this option. If this is right, then, there is some reason to think that we can make further progress when it comes to understanding such reference-shifts by exploring the possibility of such a pragmatic explanation further. It is now time to say something more positive about how the details of such an approach might be worked out.

4. Towards a pragmatic account of reference shifting

4.1 Theoretical options

As the above discussion reveals, the intended target of the recent objections raised against pragmatic approaches to reference-shifting in indexicals is a specific class of pragmatic explanations, namely those carried out in terms of Grice's theory of implicatures. Although several of the objections would, if successful, be damaging to a wider range of pragmatic explanations, it seems that these critics have been working with a rather narrow conception of pragmatics. In view of the central place that Grice's theory of implicature has occupied in the philosophical literature, this is hardly surprising. Moreover, given that very little attention has been given to the pragmatic approach in the literature, its critics should not be blamed for not having taken all of the possible options into account.

A proponent of the pragmatic approach, however, should be open to other theoretical options. Arguably, the most influential alternative to Gricean is relevance theory, as first developed by Dan Sperber and Dierdre Wilson (1986). Relevance theorists claim to offer a simpler and more cognitively realistic theory than the Gricean one. They aim to explain the data that Grice's maxims were designed to account for in terms of a single principle of relevance (suitably elaborated): "that the speaker tries to be as relevant as possible in the circumstances" (Sperber & Wilson, 1986, 381).

However, neither Gricean theory nor relevance theory strikes me as the most promising way to go in developing a pragmatic approach to the cases under consideration here. Instead, I propose an account which builds on *primary pragmatic processes* as understood by Recanati (2004), i.e. processes which—in contrast to the inferential processes appealed by Griceans and relevance theorists—do not "presuppose the prior identification of some proposition serving as input to the process" (2004, 23). In the interest of space, I will not provide any detailed comparison between these three options, but I will offer some reasons for thinking that my preferred alternative has some advantage over the others, at least in one important respect. But before I do that, I need to say more about what my proposal amounts to.

4.2 Primary pragmatic processes

In *Literal Meaning*, Recanati (2004, 17; 23) distinguishes between *primary* and *secondary* pragmatic processes. Secondary pragmatic processes are (roughly) the kind of processes appealed to in Grice's (1989) theory of conversational implicature. They are post-propositional in the sense that they require that some proposition already be considered as being expressed, they are inferential, and they are conscious in the sense that normal interpreters are aware of the

proposition expressed, what it pragmatically conveys, and the inferential relation between these two contents. In contrast, primary pragmatic processes do not require that some proposition is already considered as being expressed, they are not inferential, and normal interpreters need not be aware of the context-independent meanings of the expressions used or the primary processes itself; they are typically only aware of the output (cf. Recanati 2004, 23).

Primary pragmatic processes can be linguistically mandated (bottom-up), as is the case with *saturation*, which is required in order to get a complete proposition from sentences containing e.g. genitives, like in ‘Fred’s horse won the race’. There is no complete proposition to be grasped here absent a specification of the relation that is supposed to hold between Fred and some horse (the horse that he owns/likes /bred/betted on ...etc). However, most primary pragmatic processes are optional and context-driven (top-down). The paradigm case according to Recanati (2004, 23) is *free enrichment*, which can be illustrated by examples like ‘Mary took out the key and opened the door’, which is naturally understood as conveying that Mary opened the door with the key she took out. There is also *loosening*, which Recanati (2004, 26) describes as the converse of free enrichment. For illustration, consider ‘The ATM swallowed my credit card’, which can only be made sense of by loosening the conditions of application of ‘swallow’. Finally, there is *transfer*, which is exemplified by ‘I am parked out back’, in which either ‘parked out back’ must be understood non-literally—as a property of a car owner rather than a car—or ‘I’ must be understood as referring to the car (cf. Nunberg 1995).³³ In what follows, I will focus on transfer, which, for my purposes, seems to be the most interesting one of these three varieties of primary pragmatic process.

4.3 Activation, schemata and frames

First of all, how is transfer supposed to work? On Recanati’s model, the literal meaning of each expression is accessed first. This “triggers the activation of associatively related representations”, which then, along with the literal meaning, enter the competition for being selected as the interpretation that best fits “the broader context of discourse” (2004, 28). For instance, consider the famous example, in which the term ‘the ham sandwich’ as occurring in the following sentence is (non-literally) used to refer to the person who ordered the ham sandwich, rather than the ham sandwich itself:

(20) The ham sandwich has left without paying.

³³ Nunberg (1993, 39–40) argues for the former option, while Mount (2008, 200) argues for the latter. For present purposes, I shall, as announced in footnote 6, side with Mount on this issue with respect to the examples under consideration here.

Here is how Recanati accounts for this case:

[T]he description ‘the ham sandwich’ first receives its literal interpretation, in such a way that a representation of a ham sandwich is activated; activation then spreads to related representations, including a representation of the man who ordered a ham sandwich. *All these representations activated by the description ‘the ham sandwich’ contribute potential candidates for the status of the semantic value of the expression*; all of which are equally susceptible of going into the interpretation of the global utterance. Now, the ham sandwich orderer is a better candidate than the ham sandwich itself for the status of argument for ‘has left without paying’. It is therefore the derived, non-literal candidate which is retained, while the literal interpretation is discarded. (Recanati, 2004, 29)

Here, the alternative non-literal representation *ham-sandwich-orderer* is not derived by means of any conscious inference on the part of the interpreter. It is derived through an associative process, which is blind and mechanical, and does not involve any reflection. As Recanati puts it: “The dynamics of accessibility does everything, and no inference is required” (2004, 32).

Moreover, the process just described is local rather than global, in the sense that it operates on individual expressions rather than whole sentences, and thus it does not presuppose that the literal interpretation of the whole sentence is computed. Still, it is important to keep in mind that even though such processes are local in this sense, they are nevertheless sensitive to linguistic as well as extra-linguistic context (cf. Recanati 2004, 32). Interpretation aims (among other things) at coherence, and thus the extent to which a certain set of representations becomes accessible depends on how well they “fit” with each other and the more general features of the utterance situation.

Recanati (2004, 30; 36) suggests that the relevant notion of fit can be spelled out in terms of *schemata*. For instance, the expression ‘has left without paying’ may be taken to activate an abstract schema $\langle\alpha,\beta\rangle$, which is instantiated by its literal interpretation and the derived interpretation of ‘ham sandwich’.³⁴ As Recanati points out in the passage quoted above, the ham sandwich orderer is

³⁴ As pointed out by Rumelhart et al. (1992, 20), one need not think of the schemata as (abstract) “things” stored in memory; there need be no representational object to be identified with the schema. Instead, one could conceive of schemata as being “implicit in our knowledge” and “created by the very environment they are trying to interpret” (20). However, since this issue is not central for present purposes, I will not pursue it further here.

a “better candidate” than the ham sandwich itself for the status of argument for ‘has left without paying’ (at least given that the latter retains its literal interpretation). To say that the literal interpretation of ‘ham sandwich orderer’ and the derived interpretation of ‘ham sandwich’ instantiate an activated schema is a way of making this a bit more precise.³⁵

As already suggested, which schema gets activated depends on the broad context of discourse. For instance, if the sentence uttered is

(21) The ham sandwich is getting eaten

rather than (20), the ham sandwich itself is a better candidate argument for the predicate than the ham sandwich orderer. In other words, the literal interpretations of ‘ham sandwich’ and ‘is getting eaten’ will instantiate the abstract schema activated by the latter, and this is (part of) the explanation of why the former is selected. More generally, Recanati purports to explain our tendency to prefer coherent (i.e. schema instantiating) interpretations as follows:

Interpretational success—what brings a ‘candidate’ or potential semantic value into the actual interpretation of the utterance—to a large extent depends on the candidate’s accessibility or degree of activation. Now a schema is activated by, or accessed through, an expression whose semantic value corresponds to an aspect of the schema. The schema thus activated in turn raises the accessibility of whatever possible semantic values for other constituents of the sentence happen to fit the schema. The schema itself gains extra activation from the fact that some other constituent of the sentence has a possible interpretation which fits the schema. In such a case all ‘candidates’ or potential semantic values which fit the schema evoked by some of them mutually reinforce their accessibility and therefore increase the likelihood that they will be globally selected as part of the interpretation of the utterance. Coherent, schema-instantiating interpretations therefore tend to be selected and preferred over non-integrated or ‘loose’ interpretations. As a result, schemata drive the interpretation process. (Recanati 2004, 37)

³⁵ See Mazzone (2011) for more discussion of the notion of schema and its role in pragmatics.

One would also expect schemata, or to introduce a slightly different notion, *interpretation frames*, to be activated by extra-linguistic features of the utterance situation as well as by the presence of certain linguistic expressions.³⁶ For instance, among waiters at a restaurant, the expression ‘ham sandwich’ may be so frequently used to refer to ham sandwich orderers that this interpretation is preferred over the literal one by default, even when no incoherence would result from selecting the literal one, as in “I’ll take care of the ham sandwich”. This may be explained in terms of an interpretation frame associated with this particular kind of situation, which, just like a schema, serves to activate representations that fit certain roles in the frame.

The notion of frame appealed to here figures in e.g. linguistics and cognitive psychology, and the idea of applying it to linguistic interpretation is far from new. For instance, Charles Fillmore (1976, 25) took an important part of the comprehension process for an interpreter to consist in filling in the details of frames, which he in turn took to be associatively connected with uses of certain expressions in certain kinds of situations. A frame is to be understood as “a kind of outline figure with not necessarily all of the details filled in” (Fillmore, 1976, 29), which, after it has been activated, enhances access to further linguistic material associated with it. Thus, just like Recanati’s schemata, a frame facilitates interpretation by making more material available, and by affording structure and coherence to utterance situations as perceived by interpreters.³⁷

The basic idea here is thus that the variation in interpretation in the ‘ham sandwich’ examples and similar cases are to be explained in terms of a primary pragmatic process which is triggered and constrained by features of the linguistic as well as the non-linguistic context. In particular, these features are taken to activate frames or schemata, which in turn guide the interpreter in choosing which of the activated candidates to retain. In what follows, I will show how the cases of reference-shifting in indexicals that we have considered above can be accounted for along these lines.

4.4 Linguistic context

A first observation is that indexicals seem to follow a similar pattern to ‘ham sandwich’ when it comes to variation with linguistic context. Consider the following examples:

- (20) The ham sandwich has left without paying.
(21) The ham sandwich is getting eaten.

³⁶ For further discussion of the notion of a frame, see Fillmore (1976).

³⁷ For more on frames, see Fillmore (1982), Fillmore & Atkins (1992), and Recanati (2004, 31).

- (22) I am not here now, leave a message and I'll call you when I get back.
(23) I am here now, of course, but I won't be when you hear this, so leave a message and I'll call you when I get back.

The shift in referent of 'now' between (22) and (23) on the one hand, and the shift in denotation of 'the ham sandwich' between (20) and (21) on the other, both seem to depend on linguistic context, and in a very similar way. In both of these pairs of examples, we assume that the non-linguistic context is fixed. Nevertheless, because of the difference in linguistic context between the members of the pairs, coherence considerations pushes us towards different interpretations: the referent of 'now' is taken to be the time of playback in (22), and the time of recording in (23), while the denotation of 'the ham sandwich' is taken to be the ham sandwich orderer in (20), and the ham sandwich itself in (21).

The same kind of similarity can be observed for intra-sentential shifts, as in the following examples:

- (24) The ham sandwich left without paying for his ham sandwich.
(25) I know that I am on a purple square.

Assuming that (25) is uttered in a standard Monopoly situation, coherence considerations pushes us towards taking the referent of the first occurrence of 'I' to be the player, and the referent of the second to be his/her token. Similarly, the best candidates for filling the two relevant argument places in (24) differ across the first and second occurrence of 'ham sandwich'. In both of these cases, we can construe the process as involving the activation of a frame or schema, into which the candidate referents must be fitted, and this guides and constrains the interpreter's choice among the activated candidates.

It thus seems that everything that is needed in order for the primary pragmatic process to do its work is in place. For illustration, let us consider the Monopoly case involving (25) in some more detail. Following Recanati's (2004, 29) account of the 'ham sandwich' case, the process can be described as follows:

All occurrences of the pronoun 'I' first receive their literal interpretation, in such a way that a representation of the speaker is activated; activation then spreads to related representations, including a representation of the speaker's game piece. All these representations activated by the pronoun "I" contribute potential candidates for the status of

the semantic value of the different occurrences of the expression; all of which are equally susceptible of going into the interpretation of the global utterance. Now, the speaker herself is a better candidate than her game piece for the status of argument for ‘knows that...’, while the game piece is a better candidate for the status of the argument of ‘is on a purple square’. Therefore, the literal candidate is retained for the first occurrence, while the non-literal candidate is retained for the second.

With respect to this case, I am happy to assume that the literal referent of ‘I’ is the speaker, but in the case of the two answering machine utterances of (22) and (23), I would like to remain as neutral as possible. What I think is crucial in that case is that given that the interpreter knows how an answering machine works, both the time of production and the time of playback will be available as candidate referents for ‘now’. Moreover, given that the message in (22) is more “standard”, there will be a tendency to interpret answering machine occurrences of ‘now’ as referring to the time of playback more or less by default. However, when a different schema is activated as a result of a difference in the linguistic context, as in (23), the default interpretation will be suppressed in favour of the alternative candidate.

The important point here is that whichever of these candidates one takes to be the semantic referent, one can appeal to a primary pragmatic process in order to explain variation in interpretation across these different examples. In other words, an account of the kind sketched here can be of use to proponents of various kinds of theories, especially insofar as they are reluctant to allow speaker intentions to have a bearing on semantic content.³⁸

4.5 Non-linguistic context

The examples just considered illustrate that which schema is activated depends to a large extent on the linguistic context. But of course, the non-linguistic context is important as well, as witnessed by the difference between face-to-face or real-time phone conversations and conversations mediated by recording devices when it comes to the possibility of shifting the reference of ‘now’. Another case in point is the reference-shifting in ‘here’ between (4) and

³⁸ Of course, the extent to which one will find such a complementary account useful will depend on the extent to which the different problem cases can be handled in terms of one’s (meta)semantic theory. For instance, a proponent of the simple view will find it more useful than a proponent of e.g. Cohen’s view, since the former’s (meta)semantic theory, in contrast to the latter’s, does not have the resources to handle the answering machine puzzle.

(5), which clearly depends on the difference in non-linguistic context between the two utterances.

- (4) It is sunny and warm here today.
(5) I am not here today.

One of them is made as part of a message written on a postcard, while the other is part of a message written on a note to be put on an office door, and this seems crucial to how we are inclined to interpret them. Similarly, a crucial part of the explanation of the relatively smooth shift in interpretation between (22) and (23) is that they are understood as occurring as part of recorded messages. In contrast, suppose that Mrs Horne utters (26b) or (26c) in response to (26a) while talking on the phone with Mark:

- (26a) Will you be at the office tomorrow 9 a.m. sharp?
(26b) Yes, I will be there now.
(26c) No, I am not here now.

General coherence considerations would seem to count in favour of taking the referent of ‘now’ in (26b) or (26c) to be 9 a.m. the day after the phone call, and the referent of ‘here’ in (26c) to be the office. Moreover, this time and this place are both clearly salient in the conversation, so there is no reason to think that they are not cognitively accessible in that sense. Yet, they do not seem available as (pragmatic) referents for ‘now’ and ‘here’ in this scenario.

The explanation for this, I suggest, is to be found in the temporal distribution of the different parts of the utterance. As pointed out above, what makes two different times available as the referent of ‘now’ in the answering machine cases is that the sentences are understood as being part of pre-recorded messages, i.e. as utterances whose time of production and time of tokening comes apart. This provides the interpreter with a salient *connection* between the two, which is missing in face-to-face and real-time phone conversations. So, the idea here is that it is not enough that some other time is salient in the conversation in order for it to be genuinely available as the (pragmatic) referent of ‘now’; it must also be appropriately related to the time of production (or tokening, depending on which one we take to be the semantic referent).

Similarly, any candidate referent for ‘I’ must be appropriately related to the speaker in order to be genuinely available as a pragmatic referent. Even if Mark is the most salient person in the room, Mrs Horne cannot (in a normal

Monopoly situation) shift the reference of ‘T’ to Mark or his game piece.³⁹ The same thing goes for the ‘ham sandwich’ cases: the process leading to the transfer turns on there being an appropriate relation between the literal and the non-literal referent. For instance, the reference of ‘the ham sandwich’ cannot be shifted to the person who ordered the tuna sandwich.⁴⁰

Exactly what it takes for a relation to count as appropriate in the relevant sense seems to depend on specific features of the situation in which the utterance is made. So, the context does not only provide enabling conditions for pragmatic reference-shifting; it also provides constraints. This may be cashed out in terms of restrictions on admissible values for slots in the schema, which in turn depend on which relations counts as appropriate according to the contextually activated frame of interpretation. At this point, things are likely to get rather complicated. In particular, there seems to be no simple and general answer to the question of what appropriateness amounts to here. Of course, more needs to be said about this in the course of developing the theory, but this is an issue that I will leave for another occasion. Instead, I will end this section with some general remarks on the benefits of a pragmatic account of the kind I have sketched.

4.6 Inference and conscious availability

In view of the findings so far, there is some reason to be optimistic about the possibility of developing a pragmatic approach to the cases of reference-shifting under consideration here in terms of primary pragmatic processes. The framework sketched offers resources for spelling out both how the reference-shifts are triggered and how they are constrained. Of course, for all that has been argued so far, it may be possible to accomplish this in a Gricean or relevance theoretic framework as well. So, one may ask, why not go for one of these alternatives instead?

Well, they may be worth exploring too, and there may well turn out to be some Gricean or relevance theoretic account that can do the work. As I cannot consider all conceivable ways of working out an account of either of these kinds, I cannot exclude that possibility here. Nevertheless, from a more general theoretical perspective, I think there are some advantages of going for the kind of account I have suggested. In contrast to the account sketched above, both Gricean and relevance theoretic accounts construe pragmatic processes as inferential procedures. This is problematic as the inferential construal appears

³⁹ This might in fact be possible in cases involving overt pretence, e.g. if Mrs Horne were to utter ‘I’m on a purple square’ while imitating Mark’s voice or characteristic gestures. But this is a special case, which I will put aside for present purposes.

⁴⁰ The importance of there being a systematic relation between the literal and derived content is emphasized by Nunberg (1995).

to conflict with the first-person phenomenology of at least some of the cases at hand. For instance, it does not seem that we perform such inferences in interpreting (25).

In response to this, it might be pointed out that the inferences in question need not be something that actual language users explicitly perform. Gricean accounts are typically not even intended as a realistic model of the actual cognitive process, and as regards relevance theoretic accounts, the inference may well be taken to be unconscious in the sense that the interpreter is not aware of the fact that she is making an inference. For instance, Sperber (1997, 77–78) distinguishes such unconscious inferences—which are taken to be (relatively) effortless, fast, and spontaneous—from conscious inferences, which are taken to be voluntary, more effortful, and slow.

However, it is far from clear that this distinction really provides for a better fit between the account and our first-person phenomenology. In fact, there seems to be a relevant sense in which the inferences characterized as “unconscious” by Sperber are consciously available to the interpreter, namely the one appealed to by Recanati (2004, 42). According to Recanati, a conscious inference takes place whenever one judgement is grounded in another judgement, and both of these judgements are available—i.e. consciously accessible—to the judging subject. He emphasizes that the inferences underlying the retrieval of conversational implicatures need not be effortful, slow or under voluntary control. On the contrary, they are “typically spontaneous: the inference is drawn more or less automatically” (2004, 42). Nevertheless, they satisfy the availability condition, and thus they are conscious in Recanati’s sense. So, even if the inference that (purportedly) leads up to the preferred interpretation in the cases we have considered is neither voluntary nor explicit, it may still be conscious in Recanati’s sense.

If we think that the interpretations on which the inferences are supposed to be based are not always consciously accessible (in Recanati’s sense) to normal speakers, this will be enough to cast doubt on inferentialist accounts, at least insofar as they purport to model actual cognitive processes. So, the question is, are we always conscious of the literal contents of the cases we have considered, and the inferences that are supposed yield the derived interpretations? While I cannot settle this largely empirical question here, first-person phenomenology does seem to speak in favour of a negative answer, at least with respect to the cases under consideration here. In a normal case of interpreting e.g. (25), it just does not seem that there is any available level of consciousness at which the literal content (that the player herself knows that she herself is on a purple square) is represented, and the same holds for the purported inference from this literal content to the derived content (that the player herself knows that her game piece is on a purple square).

Whether one takes considerations like these to count against inferentialist accounts is likely to depend on one's theoretical outlook. Some may be inclined to give priority to an explanation at the level of functional description, where the goal is to show how utterance interpretation could be a rational achievement,⁴¹ while others may prefer explanations at the level of actual cognitive mechanisms. If one belongs to the former category, the fit between the theory and cognitive reality will seem relatively unimportant, but if one belongs to the latter, one need to take mismatches between the theory and the mechanisms it is supposed to model very seriously. Is there any reason to favour any of these theoretical preferences over the other? On the one hand, I think there is reason to strive for an account that fits as well as possible with cognitive reality, but on the other hand, I do not think that these explanatory projects are necessarily to be considered as rivals. One may acknowledge the legitimacy of the former project, while taking the latter to be crucial insofar as it aims to account for the actual mechanisms that underlie these achievements.⁴²

In other words, providing rational reconstructions in terms of an inferential account may be considered to be part of a respectable philosophical project, but that does not mean that we should not also try to relate these reconstructions to empirical theory. And when it comes to this latter project, accounts of the kind sketched here has an important advantage in that the associative processes governing the dynamics of activation and accessibility appealed to are fairly well documented in psychology and neuroscience.⁴³ This means that this kind of pragmatic approach to the cases discussed here has the potential of offering us an explanation of the relevant data in terms of general and empirically documented features of human cognition. I take these considerations to count significantly in favour of exploring this alternative. However, as should be clear from the preceding remarks, this does not mean that inferential views must be discarded altogether.⁴⁴

⁴¹ For instance, see Soames (2008, 460) and Wright (1987, 215).

⁴² One may, for instance (like Recanati (2007, 51)) attempt to show how “dumb” processes of activation and association can mimick “smart” inferential processes.

⁴³ For a recent discussion and references, see Mazzone (2013), e.g. the following passage: “One of the best-documented facts in psychology and neuroscience is the associative dynamic in virtue of which representations are accessible—and consequently spread activation—to each other as a function of the strengthening of the synaptic connections between them, with this strengthening being in turn a function of exposition to regular co-occurrence of stimuli (in accordance with Hebbian rule).” (107)

⁴⁴ Indeed, there may be cases of reference-shifting in which secondary inferential pragmatic processes have a central role to play. For instance, when an interpreter unfamiliar with answering machines encounters an answering occurrence of a message like (8) for the first time, she may first be baffled by the absurdity of what she takes the message to say, and then

5. Conclusion

A lot remains to be done in order to turn this sketch of a pragmatic account of the cases considered into something that would deserve to be called a theory. My main goal in this paper has merely been to put this option firmly on the table, and to show that it is an alternative worth taking seriously. More specifically, my aims have been (i) to prepare the ground for such a theory by showing that the *prima facie* challenges raised in the recent literature do not give us any principled reason to exclude the pragmatic approach from the start, and (ii) to take a few first steps in working out the details of the theory-to-be, in order to give an idea of how it would work and provide a platform for further development of it.

I have tried to remain as neutral as possible among the different “conservative” semantic accounts in the offing, in order to make the pragmatic approach available to as many kinds of theorists as possible.⁴⁵ Of course, one consequence of refraining from taking a stand with respect to such controversial issues is that one gets somewhat less specific conclusions. However, this seems like a reasonable trade-off in the present context, given that my goal in this paper is not to defend any particular semantic theory. The purpose of introducing the pragmatic approach is to bring in a new perspective on the philosophical debate on indexicals, which will hopefully prove fruitful in advancing this field of inquiry.

Finally, some of the points made in section 3—e.g. concerning the limits on the applicability of standard Gricean tests—are likely to generalise to other debates in which there is a need for evaluation of pragmatic explanations. In other words, these points may turn out to have significant implications that go beyond the most immediate concerns of this paper. My hope is that the ideas and arguments presented here will contribute to a better understanding of the phenomena discussed, and a sharpening of the criteria by which pragmatic explanations are evaluated. Such a development would surely be beneficial, not only to the debate on indexical reference, but to the philosophical discipline as a whole.

take that proposition as input to a (relevance-driven) inferential process, which eventually leads her to the intuitively correct interpretation.

⁴⁵ Neither have I explicitly excluded the (less conservative) idea that the primary pragmatic processes appealed to might be truth-conditionally intrusive, in the sense of affecting the truth-conditions of the utterances in question, rather than just the communicated content. Indeed, Recanati himself defends a view according to which primary pragmatic processes are involved in determining what is said (e.g. 2004, 17). Of course, accepting this kind of view would undermine the possibility of invoking the present account in order to defend more conservative views of semantic content.

References

- Åkerman, J. 2009. 'A plea for pragmatics'. *Synthese* 170: 155–67.
- Åkerman, J. 2010. 'Communication and indexical reference'. *Philosophical Studies* 149: 355–366.
- Åkerman, J. Forthcoming I. 'The communication desideratum and theories of indexical reference.' *Mind & Language*.
- Åkerman, J. Forthcoming II. 'Infelicitous cancellation: the explicit cancellability test for conversational implicature revisited'. *Australasian Journal of Philosophy*.
- Banfield, A. 1973. 'Narrative style and the grammar of direct and indirect speech'. *Foundations of Language* 10: 1–39.
- Bach, K. 1999. 'The myth of conventional implicature'. *Linguistics and Philosophy* 22: 327–366.
- Barker, S. 2003. 'Truth and conventional implicature'. *Mind* 112: 1–33.
- Cohen, J. 2013. 'Indexicality and the puzzle of the answering machine'. *Journal of Philosophy* 110: 5–32.
- Cohen, J. & E. Michaelson. 2013. 'Indexicality and the answering machine paradox'. *Philosophy Compass* 8: 580–592.
- Davis, W. 1998. *Implicature: Intention, Convention, and Principle in the Failure of Gricean Theory*. Cambridge: Cambridge University Press.
- Fillmore, C.J. 1976. 'Frame semantics and the nature of language'. *Annals of the New York Academy of Sciences* 280: 20–32.
- Fillmore, C.J. 1982. 'Frames and the semantics of understanding', *Quaderni di Semantica* 6: 222-54.
- Fillmore, C.J. and S. Atkins. 1992. 'Toward a frame-based lexicon: the semantics of RISK and its neighbors', in Lehrer, Adrienne and Kittay, Eva (eds.), *Frames, Fields and Contrasts: New Essays in Semantic and Lexical Organization*. Hillsdale, N.J.: Lawrence Erlbaum Associates.
- Gauker, C. 2008. 'Zero tolerance for pragmatics'. *Synthese* 165: 359–71.
- Glanzberg, M. 2007. 'Context, content, and relativism'. *Philosophical Studies* 136: 1–29.
- Gorvett, J. 2005. 'Back through the looking glass: on the relationship between intentions and indexicals'. *Philosophical Studies* 124: 295–312.
- Grice, H.P. 1989. *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Kaplan, D. 1978. 'Dthat'. *Syntax and Semantics* 9: 221–243.
- Kaplan, D. 1989. 'Demonstratives'. In J. Almog, H. Wettstein, and J. Perry (eds.) *Themes from Kaplan*. Oxford: Oxford University Press.
- King, J. 2014. 'Speaker intentions in context.' *Noûs* 48: 219–237.
- Maier, E. 2014. 'Language shifts in Free Indirect Discourse'. *Journal of Literary Semantics* 42: 143–167.

- Mazzone, M. 2011. 'Schemata and associative processes in pragmatics'. *Journal of Pragmatics* 43: 2148–2159.
- Mazzone, M. 2013. 'Attention to the speaker. The conscious assessment of utterance interpretations in working memory'. *Language & Communication* 33: 106–114.
- Michaelson, E. 2014. 'Shifty characters'. *Philosophical Studies* 167: 519–540.
- Montminy, M. 2010. 'Context and communication: A defense of intentionalism'. *Journal of Pragmatics* 42: 2910–2918.
- Mount, A. 2008. 'The impurity of "pure" indexicals'. *Philosophical Studies* 138: 193–209.
- Nunberg, G. 1993. 'Indexicality and deixis'. *Linguistics and Philosophy* 16: 1–43.
- Nunberg, G. 1995. 'Transfers of meaning'. *Journal of Semantics* 12: 109–132.
- Predelli, S. 1998a. 'I am not here now'. *Analysis* 58: 107–115.
- Predelli, S. 1998b. 'Utterance, interpretation and the logic of indexicals'. *Mind & Language* 13: 400–414.
- Predelli, S. 2011. 'I am still not here now'. *Erkenntnis* 74: 289–303.
- Recanati, F. 1995. 'Le présent épistolaire: une perspective cognitive'. *L'Information Grammaticale* 66: 38–44.
- Recanati, F. 2004. *Literal Meaning*. Cambridge: Cambridge University Press.
- Recanati, F. 2007. 'Reply to Carston'. In M.J. Frapolli (ed.), *Saying, Meaning, Referring. Essays on the Philosophy of François Recanati*. New York: Palgrave, New York.
- Recanati, F. 2010. *Truth-Conditional Pragmatics*. Oxford: Oxford University Press.
- Romdenh-Romluc, K. 2006. 'T'. *Philosophical Studies* 128: 257–283.
- Rumelhart, D.E., P. Smolensky, J.L. McClelland, and G.E. Hinton. 1986. 'Schemata and sequential thought processes in PDP models'. In J. McClelland and D.E. Rumelhart (eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, vol. 2*. Cambridge MA: MIT Press.
- Saddock, J. 1978. 'On testing for conversational implicature'. *Syntax and Semantics* 9: 281–297.
- Schlenker, P. 2004. 'Context of thought and context of utterance: a note on Free Indirect Discourse and the Historical Present'. *Mind & Language* 19: 279–304.
- Sidelle, A. 1991. 'The answering machine paradox'. *Canadian Journal of Philosophy* 21: 525–539.
- Soames, S. 2008. 'Drawing the line between meaning and implicature—and relating both to assertion'. *Noûs* 42: 440–465.
- Sperber, D. & D. Wilson. 1986. *Relevance: Communication and Cognition*. Cambridge, MA: Harvard University Press.
- Stokke, A. 2010. 'Intention sensitive semantics'. *Synthese* 175: 383–404.
- Wright, C. 1987. *Realism, Meaning and Truth*. Oxford: Blackwell.