

Indexicality and The Puzzle of the Answering Machine

It could also be fairly pointed out that part of understanding demonstratives is knowing the rules by which they adjust their reference to circumstance...¹

A number of theorists have observed that the behavior of ordinary telephone answering machines frustrate canonical views about the semantics of indexical expressions. On the one hand, the received view of the semantics for such expressions, due to Kaplan, entails that tokens of the string (1) are false whenever they occur.²

(1) I am not here now.

On the other hand, tokens of (1), recorded on millions of answering machines and voicemail systems worldwide, apparently express something true. It appears, then, that our best theory makes predictions at odds with the (quite banal) facts, and so must be amended. But how?

In this paper I'll defend a solution that (i) collapses onto Kaplan's classic semantics for non-answering-machine cases, (ii) departs from (and is superior to) Kaplan's view for answering-machine instances of (1), and (iii) extends smoothly to the variations of the answering machine cases that have arisen in the literature. Specifically, I'll defend a semantics according to which indexicals are interpreted relative to the contexts in which they are tokened. I'll compare this view — what I'll call the *context of tokening view* — to competitors, and urge that it succeeds where they fail. Moreover, I'll suggest that getting clear about the puzzle and the resources required for its solution holds important methodological lessons.³

1 The Puzzle: Setup and Preliminaries

1.1 Setup

I take it that little needs to be said to explain the pre-theoretical observation that answering-machine-produced occurrences of (1) (in non-screening uses) articulate something true. It requires more stage-setting to say why the standard semantics for indexicality makes the opposite prediction.

First, recall that, for Kaplan, the semantics of indexicals turns on the interaction of two aspects of their meaning — character and content — with context of use (and, ultimately, with circumstances of evaluation). The content of an expression can be thought of as a mapping from circumstances of evaluation to extensions (501). (For the special case of sentences, Kaplan encourages us to think of contents as “the what-is-said in a given context” (494) — something that maps from circumstances of evaluation to truth values, which are the extensions for sentences.) Kaplan’s second kind of meaning, character, is informally described as “a meaning rule,” that which is “set by linguistic conventions and, in turn, determines the content of the expression in every context,” and a “function from possible contexts to contents” (505). Kaplan captures the (universally acknowledged) context-sensitivity of indexicals by saying that indexicals have contextually-invariant characters that map to different contents as a function of the context in which they are interpreted: ‘I’ interpreted with respect to a context in which Barack Obama is the agent has one content, while the same expression interpreted with respect to a context in which David Kaplan is the agent has another.

Let us also follow Kaplan in construing a context c as a sequence of particulars including (possibly *inter alia*) a speaker c_A , a location c_P , a time c_T , and a possible world c_W ; and suppose that c is a context in which (1) occurs. Given this setup, we can reason as follows:⁴

- (i) An occurrence of ‘I’ refers to the agent of context c , viz. c_A . (axiom)
- (ii) An occurrence of ‘here’ refers to the location of c , viz., c_P . (axiom)
- (iii) An occurrence of ‘now’ refers to the time of c , viz., c_T . (axiom)
- (iv) An occurrence of ‘I am here now’ is true relative to a context c iff the agent of c , c_A , is located at the location of c , c_P , at the time of c , c_T . (from (i), (ii), (iii))
- (v) For any “proper” context c , at the time of c the agent of c is always located at the location of c . (In Kaplan’s formal system, this is expressed by requiring of LD-structures that “if $c \in \mathcal{C}$, then $\langle c_A, c_P \rangle \in \mathcal{I}_{\text{Located}}(c_T, c_W)$ ” (Kaplan, 1989, 542).)
- (vi) Occurrences of ‘I am here now’ always express truths relative to the (proper) contexts at which they occur — in Kaplan’s terminology, they are LD-valid. (from (v))
- (vii) Occurrences of ‘I am not here now’ always express falsehoods relative to the (proper) contexts at which they occur. (from (vi) and the usual semantics for negation)

It would appear, then, that the basic semantic machinery for indexicals has as a trivial consequence the denial of a banal, Moorean fact about the use of natural language (and one made manifest by ordinary 1970s consumer technology). Presumably, then, some modification of that machinery is called for. But what?

1.2 Three Preliminaries

Before commencing on a search for solutions, I pause for three preliminary points.

First, I emphasize that the project is to discern the correct *semantics* for indexical expressions. In terms of Grice's well-known distinction, our quest is for a theory of what the speaker's words meant on a given occasion, as opposed to what the speaker meant/intended by using those words on that occasion.⁵ Or, in Kripke's terms, we are after the correct theory of the semantic referents of indexicals, as opposed to their speaker referents.⁶ After all, the theory that the answering machine cases threaten is a theory about semantic reference.⁷ 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.⁸ Needless to say, disentangling the semantic from the non-semantic components of our overall intuitions about reference is, in general, a non-trivial matter. Nonetheless, we cannot adequately evaluate semantic theories for indexicality without ensuring that the data against which we test those theories are of the appropriate (*viz.*, semantic) type.⁹

A second preliminary point is that, while answering-machine instances of (1) motivate modifications to the classical Kaplanian semantics for 'here' and 'now', in my view they do not motivate an analogous modification for 'I'. I accept that, in the relevant instances of (1), 'here' and 'now' do not have as their contents parameters c_P, c_T of the context c . But I claim that the unique individual at issue in the (semantic) evaluation of all uses of (1) is the agent c_A of c : thus, 'I' does not shift its semantic referent in the way that 'here' and 'now' apparently do in answering-machine uses. (A notational variant: allow that 'I', 'here', and 'now' all shift their contents to parameters c'_A, c'_P, c'_T of some context $c' \neq c$, but insist that the agent c'_A of c' is identical to the agent c_A of c , while the location c'_P and time c'_T of c' are not identical to the location c_P and time c_T of c .) If this is correct, then it marks an interesting difference within the class of pure indexicals.¹⁰

A final preliminary point concerns the relationship between contexts and utterances. Famously, Kaplan (1989, Remark 1, 549) warns against assimilating his technical notion of sentences-in-contexts to utterances. He construes the former as idealizations abstracting away from the fact that utterances are extended in time, that distinct utterances cannot be simultaneous, and so forth. In fact, it is essential to several of the proposals considered below (ironically, including Kaplan's own, see §2) that the context relevant for the interpretation/evaluation of indexical-containing strings is tied in specific ways to the concrete actions of encoding/inscribing or decoding/tokening those strings. But this requires departing from Kaplan's more abstract, logically motivated approach in favor of a construal of contexts more directly aimed at the goal of understanding utterances of natural language. I believe this departure from Kaplan is well-motivated given our goals. Therefore, in what follows I will assume that it makes sense to speak about various contexts

(e.g., the context of inscription, the context of tokening) that would presumably have to be defined in terms of (potential or actual) utterances *qua* concrete actions.¹¹

1.3 Where From Here?

Given the way we have set up the puzzle, it is natural to cast our critical attention on two kinds of ingredients in the derivation of the objectionable result (vii). First, there are the axioms (i), (ii), and (iii) setting out the character of ‘I’, ‘here’, and ‘now’ in a way that identifies their content with elements of the context. Second, there are the standard assumptions about the metaphysics of contexts, expressed in premise (v). Together, these ingredients lead to the conclusion that the content of (1) (determined with respect to an arbitrary proper context) is false.

However, on reflection, it seems unpromising to respond to our puzzle by abandoning (v). For, in doing so, we would miss something that seems central in the case of the answering machine — namely, that *the true content expressed in answering machine instances of (1) is not a content involving just the parameters of the triple $\langle c_A, c_P, c_T \rangle$ comprising the context of utterance.*¹² Rather, it seems that the content expressed by such uses of (1) is about some *other* individual, location, and time.

This suggests that we should instead consider alterations to the other crucial ingredient in the generation of the puzzle — the axioms identifying the character of ‘I’, ‘here’, and ‘now’. This means that we can frame our problem as that of finding semantic principles that coincide with the Kaplanian analysis in the usual, non-answering machine instances of (1), but that, in a principled and extensionally correct way, deliver *different* contents for (1) — contents about a triple $\langle i, l, t \rangle$ distinct from $\langle c_A, c_P, c_T \rangle$ — in answering-machine and other relevant instances of (1). Put slightly differently, if we continue to think of contexts as Kaplanian sequences with agent, location, time, and worlds (and maybe more) as elements, then the problem is to find a semantics for indexicality on which all instances of (1) express contents about c' , such that c' coincides with c in standard non-answering-machine cases, but such that c' fails to coincide with c in answering machine cases.

In what follows I’ll consider and evaluate a range of proposals about how to specify the (by stipulation) content-determining context c' in the hope of deciding which, if any, gives the best solution to our puzzle.

2 Are Indexicals Ambiguous?

A natural response to the puzzle is to suggest that indexicals are (not only context-sensitive, but) ambiguous — that they can refer to elements of either of two distinct contexts. Indeed, in raising the puzzle, Kaplan himself appears to endorse this suggestion (though he does not incorporate the idea into the official semantics in “Demonstratives”):

If the message: “I am not here now” is recorded on a telephone answering device, it is to be assumed that the time referred to by ‘now’ is the time of playback rather than the time of recording. Donnellan has suggested that if there were typically a significant lag between our production of speech and its audition (for example, if sound traveled very slowly), our language might contain two forms of ‘now’: one for the time of production, another for the time of audition... (Kaplan, 1989, 491, note 12).

The idea that emerges (and is developed by Smith (1989)) is that the orthographic form ‘now’ is lexically ambiguous between the time parameter of the context of production/inscription (c^i) and that of the context of audition/tokening (c^t). Since tape-recordings/inscriptions can be moved between production and tokening, it is natural to say that, likewise, ‘here’ refers ambiguously to either the place of c^i or the place of c^t . In principle we might also extend this line to ‘I’, taking the latter to refer to either the agent of c^i or the agent of c^t (but see §1.2). If so, then since there are three indexical expressions in (1) (ignoring tense), each ambiguous between two readings, there should be a total of $2^3 = 8$ possible readings for (1).

Presumably the view would be that, while these readings collapse onto the single, standard interpretation (which Kaplan’s semantics guarantees to be false relative to any proper context) when $c^i = c^t$, tokens of (1) can nonetheless express truths when (because of the use of answering-machines, the slow travel of speech sounds that Donnellan suggests, or whatever) $c^i \neq c^t$. How might listeners then disambiguate tokens of (1) to arrive at a unique content when $c^i \neq c^t$? Perhaps pragmatic constraints — including, for example, the semantically guaranteed falsity of the standard interpretation — cause listeners to favor a specific alternative disambiguation.

While schematic, what we’ve said gives one strategy for understanding how the canonical Kaplanian semantics might be extended to account for the puzzle of the answering machine. Moreover, the proposal is a conservative extension of the classical machinery: it validates the official Kaplanian semantics in the special (but statistically typical) cases where $c^i = c^t$, and incorporates a simple, localized, yet powerful extension where $c^i \neq c^t$.

However, the proposal at hand depends on a postulated lexical ambiguity that, I claim, is not supported by the evidence. Specifically, application of a standard test involving the possibility of non-contradictory explicit cancellation under crossed readings tells against the presence of an ambiguity.¹³ To see this test in action, note that, because ‘bat’ is lexically ambiguous (between an animal and a piece of baseball equipment), there is a non-contradictory reading of

(2) Aloysius bought a bat but Aloysius didn’t buy a bat.

Namely, a reading of (2) on which its two occurrences of ‘bat’ receive distinct interpretations corresponding to different disambiguations does not express a contradiction. On this test, ‘here’ and ‘now’ do not appear to be lexically

ambiguous: once a context is fixed, there is no available reading of (3) or (4) (treating the relevant instances of ‘here’ and ‘now’ as indexicals) that fails to be contradictory.¹⁴

(3) Aloysius is here but Aloysius is not here.

(4) Aloysius is present now but Aloysius is not present now.

There is a further problem for the ambiguity view: it is not clear how, if target instances of (1) are ambiguous between eight candidate interpretations, listeners manage to single out one of them. Perhaps, as suggested earlier, in answering machine cases one or more candidate disambiguations (e.g., the canonical interpretation — that the agent of c^i is not located at the time of c^i) can be ruled out. (Moreover, if, as suggested in §1.2, ‘I’ always refers to a parameter of c^i , that would knock out four readings.) But this still leaves multiple live candidates for the interpretation of these instances of (1). And this seems problematic both because there’s no evidence that hearers take such instances of (1) to be ambiguous between multiple alternatives, and because it’s unclear what post-semantic constraints would tilt the balance uniquely in the favor of one of them (given that, *ex hypothesi*, semantic constraints fail to settle the issue).

Given these facts, I suggest that we need to look elsewhere for a solution to our puzzle.

3 The Context of Tokening View

In §1.1 we saw that answering machine instances of (1) appear to frustrate the classical, Kaplanian view that indexicals uniformly pick out parameters of c^i (the context of inscription). And in §2 we considered, and rejected, the view that indexicals ambiguously pick out either parameters of c^i or c^t (the context in which they are tokened). This suggests an obvious alternative: perhaps the indexicals ‘now’ and ‘here’ uniformly (unambiguously) pick out parameters of c^t .

As far as I know, the first published version of this view is due to Sidelle (1991), who explains it this way:

When one records an answering machine message, or writes ‘I’m not here now’ on a scrap of paper, one is not, at that time, (typically) making an utterance, or at least, making an assertion. One is not saying that one isn’t there when one is recording/writing. . . . One is rather arranging to make an utterance at a later time, or, if one likes, *deferring* an utterance. The genuine utterance(s) will occur when someone calls and hears the message. Since it is the deferred utterance, and not the deferring of it, which is the genuine utterance, ‘now’ refers to the time of the utterance, not the time of the recording (535).

Sidelle's ingenious proposal is that answering machines, notes, and the like are utterance-deferral technologies — technologies that permit one to produce in a context c^i an indexical-containing string that will be interpreted *not* with respect to c^i (as usual), but with respect to some distinct context c^t , in which that string will eventually be tokened. Call this the *context of tokening view*. Schematically, the context of tokening view offers just what we need — viz., a semantics on which indexicals are interpreted with respect to a context other than that in which they are inscribed/produced. Moreover, the view fits within a well-motivated conception of utterance deferral as a means by which speakers and interpreters exploit technologies to affect systematically the interpretation of strings.

The context of tokening view also inherits the advantages of the ambiguity view. Like the latter, it collapses onto Kaplan's semantics for cases where utterances are not deferred — where $c^i = c^t$ — as desired. And the context of tokening view predicts that answering-machine instances of (1) express a different content (namely, the content that the agent of c^i is not located at the location of c^t at the time of c^t).¹⁵ Therefore, like the ambiguity view, the context of tokening view is a conservative extension of Kaplan's classical semantics — one that inherits the successes of that semantics when $c^i = c^t$, but offers a principled and local extension to cases in which $c^i \neq c^t$. Moreover, it does this without the ambiguity view's objectionable postulation of lexical ambiguities.

Overall, then, there is a lot to like about the context of tokening view.

It has not, however, proven popular. On the contrary, the context of tokening view has been rejected uniformly in more recent literature (on the grounds of a series of baroque cases; see §§4–5). In fact, astonishingly, even Sidelle himself gives up the context of tokening view — in the very paper in which he proposed it! — in the face of such complications. He writes,

... consider grabbing a note from the kitchen table as one heads out to the car. Once out on the road, one reads 'I'm not here; out playing pool. Bob.' Though the location of the utterance is one's car, surely the referent of 'here' is the house (537).

In this passage Sidelle claims that the referent of the token of 'here' is the house, and not the location of c^t (viz., one's car). If he is right, then the context of tokening view must be rejected.¹⁶

To a fan of that view, it comes as a disappointment to lose Sidelle as an ally (especially so early on). But we should ask: is it true that the referent of 'here' in the case of Sidelle's note really is the house?

I propose that we should not accept Sidelle's intuition (*qua* intuition about the *semantic referent* of the token of 'here'), and instead should stick by the context of tokening view. This means holding that, in the case of the envisaged note read in the car, the token of 'here' semantically refers to the location of c^t , i.e., the car. Of course, it is also reasonable, given the setup, for us to think that the note's author (Bob) intended to communicate a content about his absence from the house (not the car); and it is reasonable, given the setup, for the note's audience to know that Bob intended to communicate a

content about his absence from the house (not the car). For these reasons, it is plausible that the content the note makes salient for Bob and his audience, the content on the basis of which those parties coordinate their actions, is about Bob's absence from the house (not the car). But we are independently committed to distinguishing speaker reference from semantic reference, and to thinking that conveyed content can fail to coincide with semantic content. And we are independently committed to thinking that the specific communicative intentions of interlocutors on particular occasions determine speaker reference/conveyed content more directly than they determine semantic reference/semantic content. Therefore, the intuitions listed above about what is intended, communicated, and the basis for the coordination of action give us no reason to give up an account on which the note's token of 'here' semantically refers to the location of c^t (the car). Moreover, we have reason to think that these intuitions are *not* revelatory of the semantics for indexicality: namely, the intuitions persist in a parallel case carried out in a language whose semantics for indexicality is stipulated to match the context of tokening view. But if so, then the occurrence of such intuitions in a case carried out in English cannot be used to show that English has a semantics different from the context of tokening view.¹⁷

The context of tokening view, then, offers a conservative and well-motivated extension of the received semantics for indexicality to answering-machine instances of (1) and other deferred utterances. While there are admittedly some intuitions about cases that appear to point against the view, there is reason (independent of the specific matters of interest) to treat these intuitions in non-semantic terms, therefore as not threatening to the view. As I see matters, then, the context of tokening view represents a promising approach to the semantics for indexicality that has substantial advantages, and that is not refuted by at least one sort of objection that has been brought against it.

Before we accept the view, however, I want to consider critically some of the available alternatives.

4 The Intended Context View

In a series of papers, Stefano Predelli argues that the lesson of the puzzle is that indexicals refer not to features of the context of inscription c^i , but rather to features of the *context of intended interpretation* — viz., the context c^{it} with respect to which the agent of c^i intends her utterance to be interpreted.¹⁸ In many ordinary cases, Predelli thinks, the speaker/agent of c^i intends her utterance to be evaluated with respect to the context of inscription itself; here, $c^{it} = c^i$, so Predelli's account agrees with Kaplan's. But in other cases, Predelli thinks, the agent will have other intended contexts in mind; and where $c^{it} \neq c^i$, the two accounts will come apart in the contents they assign. Predelli claims that this *intended context view* inherits the successes of Kaplan's semantics when the two agree, and better handles the cases where they do not.

4.1 Motivations

Predelli offers three kinds of cases to motivate the intended context view. First, he argues that his proposal better accounts for the answering-machine-produced instances of (1) that make trouble for Kaplan. Second, he argues that his own proposal extends to cases involving messages that go astray, which, he thinks, neither Kaplan's view nor the context of tokening view treats correctly. And third, he claims that his own account handles better than either Kaplan's or the context of tokening view sentences in the so-called historical present. I'll consider each of these motivations in turn.

We've already seen that the answering machine case raises trouble for Kaplan's classical semantics. Predelli thinks he can handle this case since, for him, the string (1) should be evaluated not with respect to the context c^i , but with respect to the distinct context of *intended* tokening c^{it} that the agent of c^i intends for her message. On this evaluation, the answering-machine-produced token of (1) expresses that the agent of c^{it} is not located at the location/time of c^{it} . Predelli assumes that the context c^{it} that the agent of c^i intends contains herself as its agent, the location of playback as its location, and the time of playback as its time. So the answering machine tokening of (1) is true just in case she herself is not located at the location of playback at the time of playback. Predelli thinks this is the correct truth-condition for the token.

A second motivation Predelli offers for the intended context view centers on intuitions about content shifts resulting from unanticipated vicissitudes befalling message-inscriptions. For example, he urges us to

Consider the anecdote of Jones, who expects his wife to come home at six, and writes 'I am not here now' at four, with the intention of informing Mrs. Jones that he is away from home at six — or, if you prefer, imagine that he records 'I am not here now' on a tape, expecting his wife to activate the tape-recorder upon her arrival. However, Jones's wife is late, and she only reads the message (or turns on the tape-recorder) at ten. Clearly the vicissitudes of Mrs. Jones do not affect the content of Jones's message. Intuitively, such content is to be established with respect to the time of *intended* decoding, and not with respect to the time when decoding actually took place (Predelli, 1998a, 110).¹⁹

I hope it is clear that the intuition Predelli cites here is incompatible with Kaplan's semantics, since the latter predicts that Jones's token of 'now' refers to the time of c^i — here 4:00, rather than either 6:00 or 10:00. His intuition is also at odds with the context of tokening view: the latter predicts that that token refers to the time of the (deferred) context of tokening, or 10:00. On the other hand, it is built into the case that the agent of c^i (Mr. Jones) has a determinate intention about the context c^{it} in which the note/recording is decoded by Mrs. Jones; and since the intended context view holds that the note's token of 'now' picks out the time of c^{it} , the theory exactly matches Predelli's intuition.

A final motivation offered for the intended context view concerns sentences in the so-called historical present, such as:

- (5) Napoleon brings his troops to the river; now he rests, and mounts his horse
(cf. Smith (1989); Romdenh-Romluc (2006)).

Again, this case apparently threatens Kaplan's official semantics, since if we take the occurrence of 'now' in (5) to pick out the time of the context of inscription (e.g., a time in 2011), (5) would inevitably express a falsehood. Moreover, it's hard to see how appeals to deferred utterances should help, since the case doesn't exploit a spatiotemporal gap between the context of inscription and the context of tokening, as the answering-machine or note-on-the-kitchen-table cases do; therefore, Predelli concludes that the example refutes the context of tokening view. Predelli's solution is instead to take the occurrence of 'now' in (5) to pick out the time parameter of c^{it} , i.e., the context intended by the agent of c^i . On this view, a contemporary utterance of (5) in a context c^i is true just in case Napoleon rests and mounts his horse at the time of the unique context c^{it} that the agent of c^i intends for the evaluation of her utterance.

4.2 Why Not to Appeal to Intended Contexts

I find these motivations unpersuasive.

To begin, for what it is worth, I do not share Predelli's intuition that 'now' in the case involving Jones's note (semantically) refers to the time of intended interpretation (6:00).²⁰ True, the case stipulates that the agent responsible for the inscription/recording of the string under evaluation intends for it to be tokened at 6:00. Consequently, 6:00 is plausibly the speaker referent — that particular that the speaker has a specific intention to pick out by his use of the token — of the relevant occurrence of 'now'. Moreover, that time is plausibly part of the content Mr. Jones intends to convey by his note/recording, and of the content Mrs. Jones (the addressee) takes her husband to have intended to convey with the note/recording. This, in turn, makes Mr. Jones's intentions about the context of evaluation potentially relevant to practical reasoning and consequent actions resulting from the inscription and tokening of the note/recording. What is less clear is whether this intention makes 6:00 the semantic referent — that particular picked out by the expression given the public conventions governing the language, together with facts about the utterance and the way the world is at the time — of the string's occurrence of 'now'. On the contrary, we have independent reasons for believing that a speaker's specific intentions to pick out some particular referent by her use of an expression can come apart from her general intentions to use that expression in the semantically conventional way (Kripke, 1977). And, once again, note that we would expect our specific and general intentions about the use of designators to come apart in cases involving imperfect knowledge even in languages *stipulated* to have a classical Kaplanian/context of tokening

semantics; hence, the fact that they come apart in such cases in English doesn't show that English fails to have a Kaplanian/context of tokening semantics.

What about uses of 'now' in historical present examples? I claim that these occurrences of 'now' are not indexicals at all, but bound variables. In (5) in particular, it is plausible that 'now' is bound to the time of (or the time following that of) Napoleon's bringing his horses to the river, i.e., the time of the event referred to by an immediately preceding discourse fragment.²¹ (Kaplan (1989, 489–490) was well aware that the words that interested him in "Demonstratives" have both indexical and non-indexical uses, and explicitly set the latter outside the scope of his theory.)

Thus, I contend that neither cases involving waylaid notes nor those involving the historical present are problematic for the classic Kaplanian semantics.

Still, I take it that that account is threatened by the original case of the answering machine, so a modification is needed. Should we, then, on the strength of the puzzle of the answering machine, adopt the intended context view? We should not.

A first reason for worry is that the intended context view wrongly predicts that users can change the semantic referents of their indexicals merely by changing their intentions (Corazza, et. al., 2002, 9; Romdenh-Romluc, 2006, 264). I can surely form the intention that my utterance of (1) be evaluated with respect to a context in which Barack Obama is speaking in Washington, DC in winter 2006 (what could stop me?). But my forming that intention does not make it the case that the content so expressed involves Obama, Washington, DC, or winter 2006. As Corazza, et. al., (2002, 9) complain, the intended context view results in a "Humpty Dumpty picture" according to which change of semantic referent is extremely easy. However, the semantic properties of public expressions (indexicals included) do not appear to vary as freely as the speaker's intended context for a string can.

Predelli (2002, 315) responds to this worry by invoking the semantics/pragmatics distinction. He insists that the intended context fixes the semantic referents of indexicals, and that the inclination to believe otherwise is explicable in pragmatic terms. He claims that the intuitions apparently telling against the intended context view reflect the realization that, without an unusual informational background, audiences won't be in a position to recover the intended referents of the indexicals, so won't be able to put what Predelli insists are the semantic contents of indexical-containing strings to epistemic use.

But this line of thought gets things backwards. Predelli is correct to insist that speaker referents (roughly, the individuals whom speakers believe they pick out by uses of their designating expressions) can diverge from semantic referents (roughly, the individuals whom speakers in fact pick out on the basis of the semantic properties of the expressions they use plus facts about the world). Barring omniscience, this sort of divergence seems inevitable. But the right diagnosis of these occasions is that the specific intentions of the speaker with respect to particular individuals that they should

be designata of the relevant expressions *don't* determine semantic reference. Predelli's proposal is, in effect, that the private and inaccessible intentions of speakers rather than any public features of expressions, world, and context, are determinative of semantic reference for indexicals (hence the semantic contents, truth-conditions, and truth values of the host expressions in which they fall). The problem is not (*pace* Romdenh-Romluc, 2006, 264–265) that the resulting conception of content would be epistemically unavailable to audiences — familiar Gricean considerations should make palatable the idea that the content available to ordinary audiences can diverge from semantic content *per se*. Rather, it is that the communicative intentions of speakers appear not (in general) to be constitutive of semantic content. It is possible, by forming a heterodox intention, to change the speaker referent of your designator; but forming such an intention does not, by itself, change the semantic referent of your designator. I conclude, then, that the charge of Humpty Dumptyism leveled against the intended context view stands.

What these considerations show is that a speaker's allegedly reference-constitutive intention can pick out what are, intuitively, the wrong referents of the indexicals uttered. But there are also cases where the speaker has either no intentions about the context, or has intentions that are too thin to single out a determinate context.

Consider, for example, the author of a book who writes a preface containing (6), or the marooned sailor who writes (7) in a message before sending it off in a bottle:

(6) The book you now have in your hands was a long time in coming.

(7) I am now stranded on a small island somewhere in the Pacific.

An important feature of these cases — one arguably central to the purposes of the imagined inscribers/utterers — is that the agents involved don't have specific intentions about the context of tokening beyond their hope that there should be one. If the content of the indexicals is, per the intended context view, the time/place of the context the speaker intends for the interpretation of her message, it's hard to see that the relevant instances of (6), (7), and their ilk should have semantic contents at all.²²

Taken together, I take these considerations to show that the context relevant for the interpretation of indexicals cannot be identified with the speaker's intended context.

5 The Recognized Context View

Consider, then, a proposal, due to Romdenh-Romluc, on which the context relevant to the interpretation of indexicals is that recognized/expected by the utterance's audience.²³

5.1 Motivations

To motivate this proposal, Romdenh-Romluc invites us to consider a case in which an imagined character, Simeon, mistakes a day on which he is to read the news on television for a day on which he is to recount history on television:

He goes on air, and thinking of the Norman Conquest says,

...Now the French are invading England!

Simeon intends his utterance to be interpreted with respect to 1066.

But everyone listening to his news programme identifies 2003 as the context of interpretation, thus taking 'now' to refer to 2003. They are subsequently horrified as they think that France has declared war on England (Romdenh-Romluc, 2002, 266).

Romdenh-Romluc's intuition is that, in making his utterance in c^i , Simeon says something false about 2003, rather than something true about 1066. On the negative side, this intuition is at odds with Predelli's intended context view, which takes the relevant context c' to be that intended by the agent of c^i — here, a context whose time component is 1066. On the positive side, Romdenh-Romluc's intuition is that c' is, instead, something like the context recognized by the audience of c^i — a context whose time component is 2003. However, she is alive to the concern that, just as Simeon can be misinformed, confused, incompetent, or ignorant about the context of utterance, so, too, actual audiences can be misinformed, confused, incompetent, or ignorant about the context with respect to which indexicals should be evaluated. If they are mistaken in any of these ways, she (reasonably) doesn't want to say that the context they take to be relevant fixes the content of the indexicals uttered — rather, she wants to say that such mistaken audiences *misidentify* the content of the utterances they hear. Therefore, she adverts to ideal, rather than actual, audiences, and claims that the context c' relevant to the interpretation of an indexical-containing string is that context recognized by a competent and attentive audience making use of cues she would reasonably take the speaker to be exploiting (Romdenh-Romluc, 2006, 274).

The proponent of the recognized context view might also attempt to motivate her view by appeal to the cases adduced by Predelli on behalf of the distinct intended context view. For example, recall that, in the case of Mr. Jones's waylaid note to Mrs. Jones, Predelli claimed that the note's token of 'now' picked out the time of its intended interpretation, 6:00 (rather than the time of inscription, 4:00, or the time of actual playback, 10:00). But, in the case described, 6:00 is not only the time at which the note's author intended it to be decoded; 6:00 is also the time that the note's (competent and well-informed) audience would identify as relevant for the note's decoding — the scenario stipulated that Mrs. Jones knows she was expected to arrive and decode the message at 6:00. Thus, here 6:00 is the time of a context c' that is both (i) intended by the agent of the context of utterance c^i , and (ii) identified by the audience of c^i . Therefore, one who shares Predelli's intuition that

the referent of the note's token of 'now' is 6:00 might take that intuition to support the recognized context view. Likewise, it is plausible that 'now' in a "historical present" utterance of (5) in context c^i picks out a time parameter c'_T of a context c' that is both (i) intended by the agent of c^i and (ii) identified by the audience of c^i as salient. Therefore, if we assume that the semantics of indexicality must explain why 'now' in (5) picks out a time in the early 19th century (an assumption I rejected above), then it should be counted an advantage of the recognized context view that it can deliver that result. Finally, the recognized context view offers hope of accounting for the original case of the answering machine. For if an answering-machine instance of (1) is inscribed in an utterance in a context c^i , it's not implausible that the audience parameter of c^i is the person calling the machine at a time t' later than c'_T ; and it's not implausible that that audience of c^i identifies a context c' whose time parameter is t' (rather than c^i) as relevant for the interpretation of the indexicals in (1).

5.2 Why Not to Appeal to Recognized Contexts

Despite these advantages, there are reasons to resist the recognized context view.

To begin, the view is underspecified. Grant, concessively, that the notions of competence, attentiveness, and what it is reasonable for a speaker to exploit as a cue can all be developed satisfactorily. Even so, it's not obvious that there is one unique context that would be identified by the idealized audience whose existence we are granting, or how the process of context-identification by such an idealized audience would work. This makes it hard to know just what the proposed theory predicts about cases.

More seriously, however, it seems that the view makes unrealistic assumptions about the psychological states of individuals (viz., audiences), and elevates these assumptions into preconditions for successful indexical reference.

First, Weatherson (2002) points out that there are cases in which the (idealized) audience reasonably but wrongly identifies what is intuitively the wrong context for evaluation, and so is misled into consideration of the wrong content.²⁴ One example Weatherson uses to show this involves a trick played on a colleague whose full departmental mailbox suggests that he hasn't checked his mail in several days. Suppose, on a day when a faculty meeting will take place, I leave an undated note in his mailbox reading

(8) There is a faculty meeting today.

I follow Weatherson in thinking that this note says something true about the day on which I leave it in my colleague's mailbox. But when the colleague retrieves the note on that day, he may — reasonably but wrongly — identify an earlier day as part of the context relative to which to interpret the note's occurrence of 'today'. Relative to the context identified by the competent

audience, the note says something entirely different — here, something false — about an earlier day.²⁵ Such cases, then, show that competent audiences can identify the wrong context for the interpretation of indexicals, contrary to the recognized context view.

Moreover, it seems that competent audiences can fail to identify any sufficiently determinate context for the interpretation of indexicals. For example, let us modify Weatherson's case by imagining that the department chair makes the decision every morning about whether to hold a faculty meeting on that day, and leaves a note in faculty mailboxes announcing the meeting on just those days where there will be one. Suppose there hasn't been a meeting (or a note) for weeks, but that there will be one later today, and suppose I have not checked my mailbox in several days. When I finally check my mailbox, I see the note, dutifully left by the chair, and containing a token of (8) to inform me about today's meeting. Although I am a (linguistically, if not professionally) competent and attentive audience exploiting available cues, I can't identify uniquely the context relative to which to interpret the note's occurrence of 'today' — for all I know, any context with a time component within the last several days is a viable candidate. Therefore, the recognized context view predicts that the tokening of (8) that occurs when I read the note semantically expresses no content at all. But I suggest that, on the contrary, the imagined token of (8) (semantically) expresses a content, and one that is true, despite my inability to interpret it in the imagined circumstance. Here, then, it would appear that there is a content-determining context, but that this context is not the context recognized by a competent and attentive audience as being content-determining (for there is no such context).

Perhaps the proponent of the recognized context view will reply that the idealization built into the view precludes the kinds of ignorance crucially involved in such cases. Thus, even if I or my colleague would misidentify the context controlling the interpretation of the relevant instances of (8), the theory has it that the content-constitutive context is that identified by *idealized* audiences — audiences of whom we require "competence" in a sense incompatible with making the kinds of context-identification mistakes on which the proffered counterexamples turn.

But this reply strikes me as unpersuasive. For one thing, it appears to avoid the troublesome cases only by the introduction of a label ('competence', 'idealization'). The proposal comes with no account of just what information must be available for an audience to count as sufficiently competent or idealized; and, given that there are so many possible ways in which actual audiences can be misled about contexts, it's hard to see that there could be such an account. Of course, we might stipulate that competent/ideal audiences must, by definition, have whatever information it takes to identify the unique context that determines the semantic values of the indexical expressions they hear. But since the recognized context view was offered as a theory of what that semantically relevant context was supposed to be, we cannot build into the view an appeal to this kind of knowledge without giving up the stated aim of the proposal.²⁶

Viewed from a distance, the intended context and recognized context views fail for similar reasons. On both proposals, the content of the relevant expressions depends too directly on the doxastic states of the parties to specific utterances (the agent, the audience). In this sense, both of these proposals are instances of what Kaplan calls “subjectivist semantics” for indexicality.²⁷ (Kaplan contrasts subjectivist semantics against a “consumerist” model (602), on which expressions come to individual language users with their semantic properties already more or less in place, prior to particular occasions of usage.) The problems we found for both of the proposals we’ve been reviewing are characteristic of semantic subjectivism. They show up whenever the allegedly content-constitutive subjective states in question are insufficiently accurate or specific to do the required semantic work. Thus, just as Kripke (1980, lecture II) showed that subjective states of users of proper names can single out no individual (the Jonah case), many individuals (the Feynman case), or the wrong individual (the Gödel case), we have found that subjective states of the agents or audiences of indexicals can come apart from the semantic values of these expressions in similar ways.²⁸ These reflections suggest that subjectivism fails for indexicals just as it fails for proper names: the semantic content (*a fortiori*, the content-determining context) of token indexical expressions is not fixed by the doxastic states of either the agent or the audience.^{29 30}

6 The Context of Tokening View Revisited

Let’s take stock.

We were aiming for a semantics on which indexicals pick out elements of the context of utterance in standard cases, but elements of a distinct context in answering-machine-involving and related cases. We found reasons to avoid the postulation of ambiguities, and so began the hunt for a univocal and extensionally adequate semantics that would distinguish between the cases of interest in a principled way. We next considered three univocal accounts: the context of tokening view (§3), the intended context view (§4), and the recognized context view (§5).

We found that each of these three views amounts to a conservative and non-ambiguity-involving extension of the classical Kaplanian view. Moreover, each comes with a principled explanation of what about the non-standard cases requires such an extension (that they involve utterance-deferral, or that the speaker intends her string to be interpreted with respect to a distinct context, or that an idealized hearer identifies a distinct context as relevant for interpretation). Since each of the theories appears to answer the main theoretical desiderata, it is unsurprising that the literature comparing these positions has centered on matters of extensional adequacy: the question has been which, if any, of them makes the intuitively correct predictions about a range of (increasingly bizarre) cases.

What *is* surprising (and, I claim, erroneous) is the standard assessment that the context of tokening view obviously falls short, and that the remaining issue

is to decide between the intended context and recognized context views. For, on the one hand, I have urged that, because of their subjectivist character, both the latter views deliver extensionally incorrect verdicts about ordinary cases where the knowledge of speakers/hearers falls short of omniscience (§§4–5). On the other hand, I argued that the objections about extensional adequacy on the basis of which theorists have rejected the context tokening view turn on inappropriately running together semantic and speaker reference (§3). If I am right, then, the context of tokening view is the only general, principled, and univocal semantics for indexicals that has *not* been refuted by considerations about extensional adequacy.^{31 32}

But I believe we can say more than this in support of the context of tokening view: I now want to offer a positive reason for accepting it. Briefly, the reason is this: by accepting the context of tokening view as the semantics for indexical reference, we thereby gain an explanation of the referential intentions of speakers and beliefs of hearers with respect to the interpretation of indexicals — intentions and beliefs that we have reason to think govern extragrammatic aspects of our understanding of indexical-containing utterances. Therefore, we should accept the context of tokening view because it makes available an explanation of features needed to account for the way speakers and hearers use indexical language.

This deserves elaboration.

The first point is that, as proponents of the intended context view stress, speakers typically have intentions about the contexts in which their indexical-containing utterances will be tokened, and with respect to which the indexical expressions in those utterances should be interpreted. The second point is that, as proponents of the recognized context view stress, competent and attentive audiences typically recognize particular contexts as those in which speakers intend their indexical-containing utterances to be tokened, and with respect to which the indexical expressions in those utterances should be interpreted. I have argued against treating the contexts singled out by such intentions/beliefs as fixing the semantic content of indexical expressions. But this is not to deny that such intentions and beliefs are present, or that they contribute significantly to the content that speakers/audiences communicate by means of indexical-containing utterances. Indeed, I believe that much of the attraction of the intended/recognized context views comes from noticing the powerful influence of such intentions and beliefs on our intuitions about what indexical-containing utterances communicate.³³ Thus, even if we disagree with proponents of those views in treating the relevant intentions and beliefs as not fixing the contexts that determine the semantic interpretations of indexicals, we should agree that such intentions and beliefs explain important aspects of the use of indexical language.

But this raises two questions. Why do speakers entertain intentions about the contexts in which their indexical-containing utterances will be tokened? And why do competent and attentive audiences recognize contexts as those in which speakers intend their indexical-containing utterances to be tokened? I suggest this is because *they take the condition of being the context of tokening to be*

relevant to fixing the semantic contents of indexicals. That is, the reason a speaker *S* forms the intention that her indexical-containing utterance will be tokened in *c'* is that: (i) *S* wants to communicate a certain content to her audience *A* with certain (indexical-containing) linguistic material; (ii) *S* believes that if that linguistic material is tokened in *c'*, then it will be interpreted relative to *c'*; (iii) *S* believes that if that linguistic material is interpreted relative to *c'*, then it will have as its semantic content just the content she wishes to convey by uttering it. Likewise, the reason a competent and attentive audience *A* recognizes a particular context *c'* as that in which the speaker *S* intends her indexical-containing utterance to be tokened is that: (i) *A* construes *S* as wanting to communicate a certain content with the indexical-containing linguistic material *S* did utter; (ii) *A* believes that *S* believes that if that linguistic material is tokened in *c'* then it will be interpreted relative to *c'*; (iii) *A* believes that *S* believes that if that linguistic material is interpreted relative to *c'*, then it will have as its semantic content just the content *A* believes *S* wishes to convey by uttering it.³⁴

The point, then, is that we can understand why speakers have (correct or incorrect) intentions about, and why hearers form (correct or incorrect) beliefs about speakers' (correct or incorrect) beliefs about the contexts in which their utterances are tokened given the supposition that being the context of tokening is relevant to the semantic interpretation of those indexicals. But it is hard to understand the presence of those intentional states without that supposition. Therefore, we should accept the context of tokening view, according to which being the context of tokening is indeed relevant to the semantic interpretation of indexicals.

Here is an analogy. Kripke (1977) points out that a speaker can successfully communicate information about the emotion of a teetotaler woman drinking sparkling water by the use of (10), even if the initial definite description lacks a denotation:

(10) The man over there drinking champagne is happy tonight.

On the standard story, (10) succeeds communicatively in the envisaged scenario because (i) there is a specific individual *a* about whom the speaker wishes to communicate with his words; (ii) the speaker — here, erroneously — believes that *a* satisfies (or will be efficiently identified by his interlocutor as the ostensible satisfier of) the condition of being a contextually unique man in a demonstrated location drinking champagne;³⁵ (iii) the speaker believes that his utterance of (10) will (semantically) express a content involving the condition of being a contextually unique man in a demonstrated location drinking champagne, and therefore that his utterance of (10) will convey content about *a*. But notice that this explanation of the communicative success of (10) turns on the assumption that the speaker believes *a* satisfies (or will be efficiently identified as the ostensible satisfier of) the condition of being a contextually unique man in a demonstrated location drinking champagne. And it is hard to see why the speaker would have any such belief about that particular condition unless that condition were (and were believed by the speaker to be) relevant to

the semantic interpretation of the linguistic material he used. The supposition that that condition is semantically expressed by his linguistic material explains why the condition would be (and be taken to be) relevant to the semantic interpretation of the linguistic material, and so allows an explanation of how he extrasemantically conveys what he does convey in using it.

Likewise for indexicals: the best explanation of important (but not exclusively semantic) features of our use of answering-machine involving instances of (1) and the other cases we have examined turns on speakers' intention to pick out the contexts in which they believe (rightly or wrongly) that their utterances will be tokened. The presence of that intention is explicable if, as per the context of tokening view, indexicals refer to parameters of the context of tokening, but not otherwise.³⁶

7 Conclusion

The resolution of the puzzles raised by answering machines matters because context-sensitivity is a pervasive feature of natural language, of which the behavior of pure indexicals is supposed to be a simple example. I believe that, contrary to the standard assessment, the context of tokening view offers the best answer to these puzzles. I do not claim to have shown that the context of tokening view successfully evades every objection, but I believe that there is more to be said in its favor, and less to be said in its disfavor, than commonly thought. If I am right, it may be considered a starting point upon which a more general understanding of natural language context-sensitivity might be built.³⁷

Notes

¹Davidson, D. (1967). "Truth and meaning." *Synthese*, 17, 304–323, at 319.

²Kaplan, D. (1989). Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In J. Almog, J. Perry, and H. Wettstein, editors, *Themes From Kaplan*, pages 481–563. Oxford University Press, New York. (Widely circulated in mimeograph from c. 1977.) While below I will use Kaplan's apparatus, which I take to be the received and best worked out semantics for indexicality, I believe the problems under discussion can be raised for other attempts to theorize systematically about the target expressions, such as those of Davidson (1967) or John Perry (1977). Frege on demonstratives. *The Philosophical Review*, 86, 474–97.

³My position has affinities with some of the claims about indexicality made by Egan, A. Billboards, bombs, and shotgun weddings. *Synthese*, 166, 251–279, 2009, and by Parsons, J. Assessment-contextual indexicals. *Australasian Journal of Philosophy*, 83, 1–17, 2010.

⁴The presentation here is adapted from that of Sidelle, A. The answering machine paradox. *Canadian Journal of Philosophy*, 21, 525–539, 1991.

⁵Grice, H. P. (1957). Meaning. *The Philosophical Review*, 66, 337–388; Grice, H. P. (1968). Utterer's meaning, sentence-meaning and word-meaning. *Foundations of Language*, 4, 225–242.

⁶Kripke, S. (1977). Speaker reference and semantic reference. In P. A. French, T. E. Uehling, and H. K. Wettstein, editors, *Contemporary Perspectives in the Philosophy of Language*, pages 6–27. University of Minnesota Press, Minneapolis.

⁷ Objection (pressed on me by Liza Perkins-Cohen and Alan Sidelle): perhaps the answering machine case should be treated non-semantically. Perhaps what is semantically expressed in a non-screening answering machine use of (1) is just the (false) content ascribed to it by the classical Kaplanian semantics, and the intuitions that a true content is expressed reflect some other conveyed content (e.g., that the agent who inscribed the message is not located at the place of playback at the time of playback, or that an invitation has been issued to the listener to leave a message). If so, then the the answering machine cases are not counterexamples to the received Kaplanian semantics.

Reply: Suppose I record/inscribe an answering machine message containing (1) 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 (1) in c_2 . Crucially, the (false) content expressed in c_2 does not seem detachable or cancellable in the way that we expect of extrasemantic content (Grice, H. P. (1975). Logic and conversation. In P. Cole and J. L. Morgan, editors, *Syntax and Semantics*, volume 3, pages 41–58. Academic Press, New York). The best account, then, is that (1) *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). (Cf. Michaelson, E. (2011) Shifty Characters. Ms., UCLA., 2011, pp. 7–10.)

⁸ Thus, while the examples adduced by Quentin Smith (The multiple uses of indexicals. *Synthese*, 78, 167–191, 1989) persuasively show that tokens of 'I' can have as speaker referents persons other than the speaker, or groups of speakers, I do not see that they demonstrate any variation in semantic referent, and so do not see them as damaging to the semantic theories Smith takes them to challenge.

On the other hand, saying that we want a semantic treatment of answering-machine tokens of (1) also means not being content to off-load the cases as un-theorized pragmatic side-effects, as per Stevens, G. (2009). Utterance at a distance. *Philosophical Studies*, 143(2), 213–221. Stevens contends that an answering-machine instance of (1) makes a sound, but doesn't constitute a Kaplanian utterance, in so far as Kaplan requires of utterances that they take place only at *proper* contexts — those whose agents are located at their locations at their times (Kaplan claims about

“improper” indexes failing to meet this requirement that “no such contexts could exist and thus there is no interest in evaluating the extensions of expressions with respect to them” (Kaplan, 1989, 509.) Stevens concludes that answering machine instances of (1) should not receive semantic interpretations, and consequently can’t be counterexamples to Kaplan’s view.

But this seems unsatisfactory. First, to my ear, the content conveyed by answering machine instances of (1) to the effect that some particular individual is not present in some particular location at some particular time can be literally true; if so, we need a semantic theory that assigns this content to these instances. Second, while Stevens’s suggestion may save Kaplan’s formal apparatus from the problematic case, it does so at the cost of not having an substantive explanation of what goes on in it. Whether the case falls inside or outside the bounds of semantic theory, we need a systematic way of saying what content is (semantically or non-semantically) expressed in it. Stevens has nothing to say in this regard, except to remark that the content communicated depends in some way on the literal meaning of the string (for which he accepts Kaplan’s semantics). But a candidate solution to our puzzle should tell us more than this (cf. Michaelson (2011), note 4).

⁹This will be an important methodological principle in what follows. I’ll use it negatively to rebut alleged counterexamples to which, some have urged, our semantics for indexicality should be sensitive, and therefore to undercut the motivations offered for various semantic proposals (§§4–5). I’ll also use this idea positively: I’ll argue that, by helping us to keep clear about which non-semantic communicative expectations are in effect, the principle can additionally provide clues about just what requirements are made by the semantics for indexicality (6).

¹⁰ Corazza, E., Fish, W., and Gorvett, J. (2002). Who is I? *Philosophical Studies*, **107**, 1–21, argue that ‘I’ shifts its semantic referent with the following case:

Joe is not in his office one day and Ben notices that a number of students keep approaching his door and knocking. They then stand around and look bemused for a while before leaving. Taking pity on these poor souls wasting their time, Ben decides to attach his “I am not here today” note to Joe’s door. The trick works; the students, instead of knocking and waiting, take one look at the note and then leave....

They conclude from the imagined note’s success in communicating the absence of Joe that its token of ‘I’ refers to Joe (rather than the note’s utterer/inscriber, i.e., the agent of the context of use, Ben).

But this seems unconvincing. For it should be common ground that what is communicatively conveyed by linguistic expressions often goes beyond their semantic content. And, indeed, what is conveyed in this case seems ripe for treatment by some standard story about the conveyance of extrasemantic content roughly in the tradition of Grice (1975). If so, then the communicative success of the note is no grounds for believing that its occurrence of ‘I’ refers to someone other than the agent of the context.

For parallel reasons, if I buy a recording for my own answering machine of a celebrity uttering the string (1), or if I ask or trick a friend into recording an inscription of (1) on my answering machine, the success of the recording in conveying information to callers about my location doesn’t give us reason to reject the hypothesis that the recording’s tokens of ‘I’ refer to the celebrity or friend who recorded it (*pace* Romdenh-Romluc, K. (2006). I. *Philosophical Studies*, **128**(2): 257–283, at 259, 262; and Michaelson (2011), 12ff.)

¹¹Understood in this modified way, Kaplan’s theory comes much closer to views of indexicality on which the reference of indexicals is defined relative to (potential) acts by a speaker (e.g., Davidson (1967); Burge, T. (1974). Demonstrative constructions, reference, and truth. *The Journal of Philosophy*, **71**(7), 205–223; Higginbotham, J. (1994). Priorities in the philosophy of thought, i. *Proceedings of the Aristotelian Society, Supplementary Volumes*, **68**, 85–106.; Higginbotham, J. (2002). Competence with demonstratives. *Noûs*, **36**(Issue Supplement s16), 1–16.) Thanks here to an anonymous referee.

¹²Here and elsewhere I omit the world parameter in specifying contexts.

¹³Other standard tests for lexical ambiguity are hard to apply to the case. I omit discussion for reasons of space.

¹⁴We must fix a context because there is a true/non-contradictory reading of (4) where Aloysius leaves in the middle of the utterance. (Thanks to Andy Egan.)

The parenthetical restriction to indexical uses is needed because plausibly (3) and (4) have non-contradictory readings where at least one instance of the relevant forms is read as a demonstrative rather than an indexical. But since the puzzle we are discussing concerns the semantics for pure indexicals, demonstrative uses of 'here' and 'now' are irrelevant. For the same reason, the observation that there are false utterances of (1) involving 'here', and 'now' interpreted as demonstratives (rather than pure indexicals) should not be taken as counterexamples to Kaplan's semantics (*pace* Smith (1989); Vision, G. (1985). I am here now. *Analysis*, 45(4), 198–199).

¹⁵Here I assume that 'I' does not shift its content in answering machine instances of (1); see §1.2.

¹⁶Sidelle goes on to suggest tentatively two further specific proposals for understanding why the token of 'here' refers to the house: either that the token refers to a parameter of the context intended by the speaker, or that the token refers to a parameter of the context expected by the audience. These suggestions correspond to the views of Predelli and Romdenh-Romluc discussed in §§4–5.

¹⁷Here I mirror a well-known argument of Kripke (1977, 265–266) against taking the referential uses of definite descriptions identified by Keith Donnellan (Reference and definite descriptions. *The Philosophical Review*, 75, 281–304, 1966) as evidence against a Russellian semantics for definite descriptions.

¹⁸Predelli, S. (1996). Never put off until tomorrow what you can do today. *Analysis*, 56, 85–91; Predelli, S. (1998a). I am not here now. *Analysis*, 58(2), 107–115; Predelli, S. (1998b). Utterance, interpretation and the logic of indexicals. *Mind and Language*, 13, 400–414; Predelli, S. (2002). Intentions, indexicals, and communication. *Analysis*, 62, 310–316; Predelli, S. (2005). *Contexts: Meaning, truth and the use of language*. Oxford University Press, Oxford.

¹⁹Vision (1985) and Sidelle (1991) offer earlier versions of this case.

²⁰It may be interesting to notice the evolution in the responses of at least one other commentator: Perry initially accepted Predelli's intuitions (Perry, J. (2003). Predelli's threatening note: contexts, utterances, and tokens in the philosophy of language. *Journal of Pragmatics*, 35(3), 373–387), but later disclaimed them (Perry, J. (2006). Using indexicals. In *Blackwell Guide to the Philosophy of Language*, pages 314–334. Blackwell, Malden, Massachusetts).

²¹ Cf. Corazza, E. (2004). On the alleged ambiguity of 'now' and 'here'. *Synthese*, 132(2), 289–313.) I favor the same treatment of occurrences of 'now', 'here', 'tomorrow', etc. in so-called free indirect discourse, such as occurs in

(5') Tomorrow was Monday, Monday and the beginning of another school week! (Lawrence, *Women in Love*, p. 185, London:Heinemann 1971).

If this is right, then such cases are similarly ineffective as motivations for the intended context view.

²²Predelli (1998a, 114) raises a related issue with respect to the answering machine that plays back instances of (1) on many occasions, hence in many contexts, and concludes that the string can express a plurality of contents — one corresponding to each context in which it is tokened. This proposal has some plausibility, but is in tension with Predelli's official view that the reference of

the indexicals in each particular utterance of (1) is fixed by the utterer's intended context for (1). If the indexicals in the string have multiple referents, or if the multiple particular utterances of the indexicals in the string each have a referent, this suggests strongly that these referents are not fixed by the utterer's (here very thin) intentions about contexts of tokening.

²³Romdenh-Romluc, K. (2002). Now the french are invading England! *Analysis*, 62(1); Romdenh-Romluc (2006).

²⁴Weatherson, B. (2002). Misleading indexicals. *Analysis*, 62(4), 308–310.

²⁵In Weatherson's other sort of example, the author of a message says something false, but misleads the competent audience into interpreting the message with respect to a context so that its content is true. I agree with Weatherson that these cases, too, are damaging to the recognized context view.

²⁶Moreover, if idealization were a legitimate strategy for avoiding cases in which audiences fail to identify the semantically relevant context, then defenders of the intended context view should be able to invoke an analogous idealization to preclude the speaker errors in the Simeon cases that Romdenh-Romluc uses to motivate moving from the intended context view to the recognized context view.

²⁷Kaplan, D. (1989b). Afterthoughts. In J. Almog, J. Perry, and H. Wettstein, editors, *Themes From Kaplan*, chapter 18, pages 565–614. Oxford University Press, New York, at 600.

²⁸Kripke, S. (1980). *Naming and Necessity*. Harvard University Press, Cambridge, Massachusetts.

²⁹Similar considerations lead Corazza, et. al., (2002) to favor a non-subjectivist semantics for indexicals according to which it is *convention* rather than subjective doxastic states that fix the content of indexicals:

Our proposal is that, for any use of the personal indexical, the contextual parameter of the agent is *conventionally given* — given by the *social or conventional setting* in which the utterance takes place. For instance, with “now”, the setting or context in which it is used changes the time that the term refers to: if “now” is heard on an answering machine, we take the relevant time to be the time at which it is heard, and we arrive at the referent accordingly. In contrast, if we read “now” on a postcard (“the weather is beautiful now”) . . . we take the relevant time to be the time at which the words were written. Hence we get a different referent in each case (11).

I find it hard to see this as a successful answer to our puzzles. For one thing, it is difficult to assess the proposal without some explanation of how conventions “give” the content-determining context. What are the relevant conventions? How do they operate? Corazza, et. al., don't even attempt to articulate general principles to answer to these questions, but instead content themselves with telling us which person/time/place they take to be the referents of a few individual uses of indexicals, and claiming that operative conventions happen to select just those persons/times/places in those cases. But if so, then we should hope for a systematic theory of how they do, as opposed to a list of cases. (At one point (note 10) these authors compare their picture to the playing of Wittgensteinian language games; the difficulty I'm now pointing to is that, without an articulation of the rules governing such language games, it's hard to see the resulting proposal as a semantic — or any other kind of — theory.)

³⁰The force of this type of reflection seems to me to vary between classes of expressions. While I follow Kripke (and many others) in thinking that such reflections are effective in undercutting subjectivism about the semantics for proper names, I think they are much less effective in undercutting subjectivism about the semantics for pure demonstratives. From this perspective, we can think of the present dispute about the viability of subjectivist semantics for indexicality

partly in terms of how to situate the semantics for indexicals with respect to the models of proper names on the one hand and pure demonstratives on the other.

³¹For reasons discussed in §1.1, I hold that 'I' picks out the agent of the context of utterance. Thus, on my preferred version of the context of tokening view, the single context semantically controlling the interpretation of indexicals is neither $\langle c_A^i, c_P^i, c_T^i \rangle$ nor $\langle c_A^t, c_P^t, c_T^t \rangle$, but $\langle c_A^i, c_P^t, c_T^t \rangle$. (I ignore this wrinkle in the main text.)

Michaelson (2011) offers two criticisms of this claim (although he correctly notes that it is detachable from the main thrust of the context of tokening view).

First, he worries that the context $\langle c_A^i, c_P^t, c_T^t \rangle$ is objectionably gerrymandered because it contains a speaker who is not present (and is not speaking) at the time, place, and world of utterance (note 15). I agree with Michaelson that such contexts violate expectations generated by exclusive consideration of non-deferred utterances. But I find it predictable and unproblematic that such expectations fail once we allow for utterance-deferral. One can accommodate the new possibilities introduced thereby either by allowing for standard characters operating on non-standard contexts (as I propose), or by allowing for non-standard characters operating on standard contexts (as Michaelson prefers). That a theory departs in one of these ways from the classical Kaplanian view should not warrant its rejection.

Second, Michaelson (15) objects that a competent speaker of English who has never before encountered an answering machine will (despite her linguistic competence), on first hearing a token of (1), take its occurrence of 'I' to pick out c_A^t (rather than c_A^i). But I claim that the proponent of the context of tokening view can describe this case adequately. She should say that the envisaged speaker/hearer will (by virtue of his semantic competence) possess the character rule that maps the token of 'I' onto a parameter of c^i , but that his ignorance about the technologically enabled possibility of utterance-deferral may lead him to misconstrue c^t (the one context in which he interacts with the token) as the context of inscription, and so to end up with false beliefs about what contents were expressed. (Analogy: if H hears S utter 'I am tall', then H 's competence with the character-rule that 'I' refers to the agent of the context of utterance won't protect him from arriving at false beliefs about what content was expressed by the utterance if he mistakes the context of utterance for one in which S' was the agent.)

³²Michaelson (2011, 14–15) lodges two further objections against the context of tokening view.

First, he complains that the view falsely predicts that inscriptions of (9) received on postcards from vacationing friends attribute pulchritude to the place/time of tokening rather than the place/time of the postcard's inscription.

(9) It's beautiful here now.

In response, I propose that the "indexicals" in (9) are either anaphora/bound variables, bound to the place/time made salient by the picture, place-name, and date on the postcard (cf. the treatment of free indirect discourse occurrences in note 21), or demonstratives, for which the postcard itself serves as a completing demonstration. On either option, the context of tokening view avoids making the prediction Michaelson finds objectionable.

Second, Michaelson rejects the context of tokening view's prediction that tokens of (1) recorded (knowingly or unknowingly) by a friend on my answering machine say something about the friend (rather than, as per Michaelson's intuition, about me). As indicated above (note 10), I am prepared to accept this prediction and to explain away intuitions to the contrary in non-semantic terms.

³³Recall that a central objection against the subjectivist semantic proposals we considered is that, on those views, semantic values can't come apart from the subjective intentional states of speakers/audiences even when they lack relevant information, hence can't allow for semantic error in such cases. This problem won't arise for a view that couples a consumerist semantics (such

as the context of tokening view) with an appeal to the intentional states of speakers/audiences to explain what is communicated.

Moreover, because it denies that the intentional states of speakers and hearers fix semantic content, there is considerably less pressure on this kind of view to choose between the intentional states of speakers (as per the intended context view) and hearers (as per the recognized context view). On the contrary, we can now say that both sorts of intentional states fix contexts, and that the contents determined relative to those contexts interact (however that works) to result in conveyed content (which can be less than fully determinate). (The reasons given for rejecting ambiguous or less than determinate semantic contents discussed in §2 are not reasons for rejecting ambiguous or less than determinate conveyed contents.)

³⁴Obviously this kind of rational reconstruction can look unrealistically over-intellectualized. I am not claiming that speakers/audiences explicitly run through such reasoning in the course of real-time linguistic usage, but that it provides a rationalizing account of the intentional states apparently at work in speakers'/audiences' use and understanding of indexical-containing utterances.

³⁵The parenthetical qualification is needed to account for cases in which the speaker does *not* believe that the individual he wishes to talk about satisfies the condition semantically expressed by the linguistic material he uses — e.g., where the speaker uses 'the king' to talk about an individual whom both speaker and hearer believe to be a usurper, and where each knows the other believes this as well (Donnellan, 1966, 290–291). (Thanks here to Rick Grush and Mario Gómez-Torrente.)

³⁶The analogy to Donnellan's case may suggest that this line of explanation could go through even if indexicals did not semantically refer to parameters of the context of tokening, but merely brought such parameters to mind in some extragrammatic way (cf. note 35); however, this possibility is hard to square with the apparent nondetachability/noncancelability of the association (see note 7).

³⁷Thanks to Dave Barner, Craig Callender, Ivano Caponigro, Andy Egan, Mario Gómez-Torrente, Rick Grush, Eliot Michaelson, Liza Perkins-Cohen, Sam Rickless, Alan Sidelle, and Chris Wüthrich, to the editors of this journal, and to audiences at the London Institute of Philosophy, the University of Leeds, Arché, the University of Texas, Austin, and Arizona State University.