

Shifty Characters

Eliot Michaelson

University of California, Los Angeles

Abstract

In “Demonstratives,” David Kaplan introduced a simple and remarkably robust semantics for indexicals. Unfortunately, Kaplan’s semantics is open to a number of apparent counterexamples, many of which involve recording devices. The classic case is the sentence “I am not here now” as recorded and played back on an answering machine. In this essay, I argue that the best way to accommodate these data is to conceive of recording technologies as introducing special, non-basic sorts of contexts, accompanied by non-basic conventions governing the use of indexicals in those contexts. The idea is that recording devices allow us to use indexicals in new and innovative ways to coordinate on objects. And, given sufficient regularity in the use of indexicals on such devices, linguistic conventions will, over time, come to reflect this innovation. I consider several alternatives to this ‘character-shifting’ theory, but none is able to account for the data as well as the present proposal. Many face additional theoretical difficulties as well. I conclude by explaining how the character-shifting theory not only retains many of the virtues of Kaplan’s original semantics, but also coheres with a plausible view on the nature of semantic theorizing more generally.

Keywords: Indexicals, Semantics, Pragmatics, Reference

1 Introduction

In his (1989b), David Kaplan introduced what has since become the standard semantic account of indexicals (terms like ‘I’, ‘here’, and ‘now’). The theory offers an appealing

Penultimate version. Final version to appear in *Philosophical Studies*. The final publication is available at <http://link.springer.com>. DOI: 10.1007/s11098-013-0109-7.

package: it is simple, it accurately predicts our truth- and falsity-judgments for a wide range of indexical-containing utterances, and it naturally accounts for some apparent truisms regarding how indexical-containing utterances relate to the world.

Unfortunately, Kaplan’s semantics for indexicals also makes some seemingly incorrect predictions in a constrained, but still quite common, class of cases: those involving non-face-to-face communication. The problem can be illustrated by pointing out that Kaplan’s semantics predicts that (1) is false whenever it is tokened.¹

(1) I am not here now.

But there are contexts in which a token of (1) will commonly be judged true — the classic example being when it is played back on an answering machine.²

Kaplan himself recognized the potential problem posed by such cases and suggested that they should be treated as ‘special uses’ of the indexicals, but this is where he left the matter (Kaplan 1989b, p. 491 fn. 12). Best I can tell, there are two main strategies available for dealing with non-face-to-face uses of the indexicals: first, we might take our initial truth-judgments at face value and modify Kaplan’s semantics to account for them. Call this the ‘semantic strategy’. Second, we might dispute these truth-judgments and seek to explain why we routinely make such judgments in spite of their falsehood. Call this the ‘pragmatic strategy’.

Both of these strategies subsume a diverse range of specific proposals. The semantic strategy has been significantly better explored in the literature, and my own preferred response is a variant of that strategy. The pragmatic strategy, on the other hand, has been

¹Kaplan himself speaks not of tokens, but rather of evaluating ‘sentences in contexts’. Kaplan focuses primarily on utterances of sentences, but also mentions both written sentences and sentences recorded on tape. Since our focus here will be on voice recordings and playback of such recordings, it will help to have a more flexible term than ‘utterance’ to indicate of the sort of event that associates a sentence with a context and thereby makes it ripe for semantic evaluation at that context. I opt for the term ‘token’ (and the associated event ‘tokening’), which I mean to cover utterances, inscriptions, and also playback of recordings — all of which, I take it, can appropriately be assigned a meaning and truth-value relative to a context.

²This observation has been pressed as an objection to Kaplan’s proposed semantics for some time. The earliest instances of this challenge are Vision (1985), Smith (1989), and Sidelle (1991).

largely neglected. While I do not endorse the pragmatic strategy, one of the goals of the present inquiry will be to rectify this oversight and put at least the most obvious variant of that strategy firmly on the table.

The miscellany of semantic strategies can be grouped in several ways. One helpful division for our purposes is this: some theorists have attempted to depart only minimally from Kaplan’s original semantics, while others have argued for a more radical break. I will canvass what I take to be the leading examples from each of these poles: with Cohen (2013) representing the conservatives and Predelli (1996, 1998a, 1998b, 2002, 2005) representing the radicals. My own proposal represents an attempt to carve out a moderate position. The idea will be to try to account for the totality of data better than Cohen is able to while avoiding the overgeneration and self-undermining worries that dog Predelli’s theory.

My basic claim will be that certain sorts of recording technologies force us to shift one aspect of the meaning of indexicals: namely, the rules that associate each indexical with a referent, relative to a context. The idea is that, when we record indexicals on certain sorts of devices, we use those terms to coordinate with our audience according to a different set of rules than those employed in face-to-face communication. This means that the rules governing the reference of ‘I’, ‘here’, and ‘now’ *change* when we begin to use certain types of recording technology — although the old and new rules are likely to agree on what these terms refer to every time they are uttered and not used on a recording.

Here is the plan for the rest of the paper: in §2, I more fully introduce Kaplan’s semantics for indexicals and the desiderata to which it is responsive. This will help to elucidate the challenge posed to that theory by answering machine cases. In §§3–4, I jump the gun a bit and introduce my own ‘character-shifting’ account of recorded indexicals. This will set up a clear contrast case as we explore, first, Cohen’s modification of Kaplan’s semantics (§5) and then Predelli’s alternative semantics for indexicals (§6). Finally, I discuss the pragmatic strategy and argue that it faces several significant *prima facie* challenges (§7). I conclude by

considering how semantics ought to relate to the historical evolution of linguistic practices. I suggest that my character-shifting approach offers a reasonable, low-cost way of responding to several desiderata in this vicinity (§8).

2 Kaplan Semantics and Answering Machines

Kaplan proposes that, relative to a context of utterance or inscription, ‘I’ refers to the speaker or writer, ‘here’ refers to the place, and ‘now’ refers to the time (Kaplan 1989b, p. 505). Hence (2) is true iff I have the property of hungriness at 4:45pm on August 16th, 2011.

(2) I am hungry now.

This is because ‘I’ refers to me, the writer, ‘now’ refers to the time of the inscription (4:45pm on August 16th, 2011), and the rest of the sentence attributes to the subject (me) the property of being hungry at the time when I wrote this sentence. That these are the correct truth-conditions to associate with this inscription of (2) will hopefully strike the reader as plausible.

In Kaplan’s semantics, the rules that associate indexicals with referents are called ‘characters’. Contexts are assumed to include, at the very least, agents, times, and places (or, more precisely, one of each) — so characters can be represented as functions from contexts to particular features of those contexts (with ‘being the speaker or writer’ serving as Kaplan’s informal gloss on what it means to be the agent of a context). The character of ‘I’ is a function from the context to its agent, ‘here’ a function from the context to its location, and ‘now’ a function from the context to its time.³ Characters, modeled as functions from con-

³Perry (2006, 2009) and Sherman (2013) go one step further in clarifying Kaplan’s apparatus by associating each indexical with a particular *contextual role* — such as THE SPEAKER or THE TIME. Such roles, as I understand them, aren’t meant to supersede character-rules. Rather, they help explain how indexicals are used by speakers to coordinate with listeners on particular objects: by guiding their attention to whatever object satisfies a particular role at the context. While I am quite sympathetic to this amendment to (or, perhaps, clarification of) Kaplan’s semantics, it introduces further complications that are best avoided in the

texts to particular features of those contexts, thus provide a stable, context-invariant aspect of the meaning of indexicals. To be clear, however, characters are not themselves what uses of indexicals denote relative to a context. Rather, indexicals are devices for picking out one or another object — and their semantic content relative to a context is exhausted by that object. In Kaplan’s terms, indexicals are thus ‘directly referential’ (Kaplan 1989b, p. 492).

Importantly, in Kaplan’s semantics, while indexicals are highly context-sensitive, they are insensitive to the intentions of the speaker. That is, indexicals pick out objects/times/places as designated by their character *independently* of what speakers hope to pick out with particular uses of indexicals, and regardless of what those speakers ultimately hope to communicate by means of their utterances. This means that the feature of the context to which a particular indexical refers cannot itself be fixed by appeal to the speaker’s intentions. This puts indexicals in sharp contrast with demonstratives (terms like ‘this’ or ‘that’), the reference of which is standardly taken to depend, at least in part, on the speaker’s mental state when using such terms.⁴

So far, I hope to have shown Kaplan’s semantics to be both simple and appealing. Now I turn to the generalizations about how indexical-containing utterances relate to the world that Kaplan’s theory was designed to capture. Consider, for instance, an utterance of (3) at Espresso Profeta at 4:48pm on August 16th, 2011.

(3) I am here now.

In this context, (3) is true iff I am in fact located at Espresso Profeta at 4:48pm on August 16th, 2011. But (3), considered as a sentence-type, demonstrates another interesting property: so long as we consider only face-to-face utterances, it is true whenever it is tokened.

This is because, any time (3) is uttered, there must be a speaker — and that speaker must

main text. If, however, the reader finds it useful to think in terms of contextual roles and satisfiers rather than in terms of functions from contexts to features of those contexts, I harbor no objection.

⁴For intention-dependent theories of demonstratives, see Kaplan (1989a), Bach (1992), and Siegel (2002). For mental-state free accounts of demonstratives, see Quine (1968) and Kaplan (1989b).

be located in a particular region of space-time. Since (3) merely attributes to that speaker the property of being located wherever and whenever she happens to be located when she utters (3), an utterance of this sentence should never fail to be true.

Kaplan accommodates this generalization about how the sentence-form (3) relates to the world by stipulating that the only contexts (represented as agent, place, time, world four-tuples) admissible in his semantics are those which contain a speaker who is actually located at the place, time, and world of the utterance (Kaplan 1989b, p. 509). Kaplan calls these ‘proper contexts’. ‘Improper contexts’, on the other hand, violate this constraint by allowing absent objects to be designated as the agents of those contexts. Considering only proper contexts ensures that tokens of the sentence-form (3) are true whenever they are evaluated within Kaplan’s semantic framework — leading him to call this sentence-form a ‘logical truth’ (Kaplan 1989b, p. 509).⁵ Likewise, the negation of (3), or (1), turns out false whenever it is evaluated on Kaplan’s semantics. This leads to its characterization as a ‘logical falsehood’.

It is this very restriction to proper contexts, coupled with the rest of Kaplan’s semantic machinery, that transforms recorded tokens of (1) into problem cases for the theory. The claim that (1) is a logical falsehood looks well-motivated so long as we restrict our consideration to face-to-face contexts. But as soon as we start considering playback tokens of (1) — which certainly seem to be well-formed tokens of natural language — we are forced to reckon with the intuition that such tokens can in fact (in the right sort of situation) be true.

⁵ Note that Kaplan intends for logical truth to come apart from necessary truth. In contrast to (1), “I might not have been here now” looks to be true on most, if not all, occasions of use (Kaplan 1989b, p. 509). Kaplan accommodates this observation by distinguishing between *contexts* — relative to which indexical-reference is determined, and which are required to be proper — and *circumstances* — relative to which modal claims are evaluated, and which are allowed to be improper. Since indexical-reference is determined relative to contexts only, not circumstances, Kaplan’s semantics is able to allow simultaneously that: (i) unembedded tokens of (1) will always be false and unembedded tokens of (3) will always be true, and (ii) tokens of (3) embedded under modals may, once the reference of their indexicals is fixed relative to the context, turn out to be false relative to a circumstance of evaluation (or, more precisely, a proposition derived from a token of (3), with the reference of each of its referring terms fixed relative to the context, may turn out to be false relative to one or more circumstance).

This conundrum has sometimes been called ‘the answering machine paradox’ (Sidelle 1991, p. 526). I’ll settle for the more plebeian ‘answering machine problem’.

The answering machine problem forces us to make a choice: we can either take our truth-judgments at face value and attempt to modify semantic theory in light of them, or we can preserve Kaplan’s semantics and reject the initial truth-judgments. As noted at the outset, I advocate going the former route. That said, I grant that the latter route has some appeal; in particular, it might allow us to preserve Kaplan’s notion of logical truth and falsity. I turn now to the task of demonstrating that semantic solutions to the answering machine problem can preserve a close cousin of logical truth and falsity with regard to sentence-forms like (1) and (3).

3 Character-Shifting

Here is the basic picture I advocate: indexicals remain context-sensitive and intention-insensitive, as Kaplan claimed. However, in contrast to Kaplan, I propose that we type contexts and allow character-rules to vary relative to context-type. In one sense, the character-rules associated with each indexical remain constant: ‘I’ returns the agent, ‘here’ returns the location, and ‘now’ returns the time of the context. However, what counts as an agent, location, or time is allowed to vary with the type of context under consideration. So, the formal notion of character — functions from a regimented set of features of contexts to objects — remains unchanged, but the informal characterization of characters as broader rules of reference is now allowed to vary relative to the context-type.⁶

In slightly more detail: in face-to-face contexts, ‘I’ refers to the speaker. However, in

⁶If the reader prefers, she is welcome to think of the proposal purely in terms of the question of what constitutes the features of agency, location, and temporality for contexts. The claim is that we find both variation in such features across contexts, but also invariance relative to particular sorts of context-types — regimented according to the type of recording technology employed at that context. In the main text, I will employ the notion of ‘character’ in the broader, informal sense — as referring to rules of reference rather than functions from features of contexts to objects — unless otherwise specified.

contexts where answering-machines are triggered and play back their messages — in which, plausibly, there is no speaker — ‘I’ refers to the owner of the line.⁷ Likewise, ‘here’ and ‘now’ in face-to-face contexts refer to the place and time of production, respectively; in answering machine contexts, on the other hand, ‘here’ and ‘now’ refer to the place and time of playback. It is thus not the formal rules associating indexicals with a particular feature of a context that shift, relative to a context-type, but rather the features associated with being an agent, time, or location. In contexts involving certain sorts of recording technologies, agency, location, and time are instantiated by different sorts of things than they are in face-to-face contexts.⁸

Several clarifications regarding the character-shifting view are in order: first, note that the character-shifting account evaluates indexical reference relative to the context of interpretation, not the context of production. By focusing primarily on face-to-face contexts, Kaplan avoided having to specify which of these contexts was relevant for the evaluation of indexical reference (in spoken contexts, the context of production usually *just is* the context of interpretation).⁹ Since it is designed to deal with recordings, for which this identity does not hold, the character-shifting theory of indexical reference needs to make a commitment in this regard. The view opts for the context of interpretation primarily because such contexts

⁷‘Ownership’ might not be quite the right notion here. As Jessica Pepp has pointed out to me, we don’t want to say that a token of ‘I’ on a work-related voicemail system refers to the company at which one is employed. ‘Being the original speaker’ is another obvious option, as is ‘being the primary user of the line’. Both of these options face their own sets of difficulties, however. I make no claim here to have found a perfect characterization of the character-rule associated with answering machine tokens of ‘I’, just to have developed a theoretical picture that allows us to productively consider this question. That said, I believe something in the vicinity of ‘ownership’ — along the lines of ‘being the recognized possessor of’ — might do the job.

⁸Context-typing can be implemented formally in several different ways. I tend to favor modifying our representation of contexts from agent, place, time, world four-tuples to agent, place, time, world, type five-tuples (with the contextual roles above specifying what it means to be the ‘agent’ or ‘time’ relative to a particular context-type). My character-shifting theory would thus claim that, quantifying over all contexts of a particular type, what constitutes an agent, time, or place remains constant.

⁹Note, however, that when the context of production and interpretation do come apart, it is clear that spoken indexicals refer to the place and time of production rather than interpretation (consider, for instance, the truth-conditions of (3) shouted over a long distance, or from a location where the speaker is not visible to her auditor). One might account for this observation either by making spoken indexicals sensitive to the context of utterance and preserving Kaplan’s original thought that ‘here’ and ‘now’ refer to the place and time of the context, or by making indexicals sensitive to the context of interpretation and specifying that ‘here’ and ‘now’ refer instead to the place and time of production. I opt for the latter route.

subsume contexts of production (given that they antedate them), but not vice-versa. This grants the view the flexibility required to make accurate predictions.

Such accuracy, however, requires some departures from Kaplan's original character-rules. Consider, for instance, a case in which my friend Tarek mails me a postcard on which he has inscribed (4).

(4) It is beautiful here now.

Kaplan's semantics predicts that 'here' and 'now' refer to the place and time of the context of inscription — thereby predicting, correctly, that (4) is true iff it was beautiful when and where Tarek wrote the postcard. My character-shifting theory, on the other hand, claims that (4) should be interpreted relative to the context of interpretation. This would leave 'here' and 'now' referring to the wrong place and time if we were to maintain Kaplan's original character-rules. However, as I above specified, the character-shifting theory actually posits that 'here' refers to the location of production, not interpretation. 'Now' was similarly modified, and 'I' can be slightly modified so as to refer to the author, rather than the speaker, when dealing with inscriptions like postcards. With these modifications in hand, the character-shifting theory is thus able to offer the correct predictions with respect to (4).

Second, the character-shifting view is intended to be relatively conservative. In particular, and in contrast to several related views (e.g. Smith (1989), Corazza et al. (2002)), it is *not* intended to capture all instances where indexicals are used to effectively communicate about objects distinct from their classical Kaplanian referents. Smith (1989), for instance, claims that 'now' can in fact be used to refer to times in the past — his example being a lecturer who pretends to be speaking from the perspective of someone at the time about which she is lecturing (pp. 170–71). I, on the other hand, would prefer to account for such pretense cases pragmatically. Similarly, Corazza et al. (2002) claim that complex, intention-insensitive conventions governing the use of indexicals will prove sufficient to ex-

plain purportedly ‘demonstrative’ uses of indexicals like ‘here’ — to refer, for instance, to particular places via maps (pp. 14–15). In contrast, I find it highly unlikely that conventions alone will suffice to determine the reference of such demonstrative uses. The problem is that, in contrast to indexical uses of ‘here’, we seem to have ample reason to posit that demonstrative uses of that term are sensitive to the speaker’s intentions.¹⁰

Third, my character-shifting account is fully compatible with the claim that indexical reference is conventional, in some fairly intuitive sense of that notion. According to standard accounts of convention (e.g. Lewis (1969)), conventions arise in response to ‘repeated coordination problems’, or situations in which a desired outcome depends on multiple agents being able to coordinate their actions over time and despite no particular action seeming *obvious* to all the relevant agents. That is, coordination problems arise where there are multiple, equally good ways of carrying out some coordinated action, but where success depends on everyone repeatedly choosing *the same way*. The character-shifting theory conceives of the use of indexicals on recording devices along just these lines: when those devices are initially introduced, there will be multiple ways that indexicals might be used on these devices. However, once a regularity in use arises, speakers will have every incentive to follow this pattern when using indexicals on this sort of recording device; deviation from such conventions will substantially decrease one’s chance of conveying a particular object-, place-, or

¹⁰The case for such intention-sensitivity is two-fold: first, not all uses of ‘here’ uttered in front of a map, even with a pointing gesture accompanying them, refer to a place represented by a point on that map. Rather, speakers can both intend to and succeed in referring to the place at which they are uttering a sentence containing ‘here’ even when this utterance is accompanied by a pointing that seems to be directed at a nearby map (consider the case of a history lecturer with a nervous tick that makes her point erratically, and who is also standing in front of a large map). It is not clear how intention-insensitive accounts, like that of Corazza et al. (2002), will be able to account for this phenomenon. Second, two tokens of ‘here’ uttered while pointing at the same place on the same map can, in fact, refer to distinct places. Consider, for instance, a flight instructor who taps a single point on a map twice and says “We’ll fly here, enter a holding pattern, and then touch down here.” Plausibly, the first ‘here’, but not the second, refers to a region some distance above the ground. Appealing to the flight instructor’s intentions offers an obvious way to fix each of these referents. In contrast, fixing reference here via an appeal to intention-insensitive conventions, as Corazza et al. propose, looks to be a difficult (if not impossible) task. Unless our conventions on map-use include appeals to discourse-context, it is highly unclear how two tokens of ‘here’ accompanied by two qualitatively identical pointings and tappings might end up referring to two very different regions of space.

time-dependent proposition by means of an indexical-containing sentence-form.

Crucially, an analogous situation does not obtain with regard to uses of indexicals that depart from Kaplanian orthodoxy but which do not involve specific types of recording technology. The problem is that no single cue can be isolated that might be used to characterize a coordination problem. Consider, for instance, uses of ‘now’ to communicate something about a past time. The stage for such uses can be set explicitly by the preceding discourse, or by making certain gestures, changing one’s tone of voice, putting on a costume, etc. The human ability to interpret such uses consistently across such a wide variety of situations is not being brought into question here; rather, what I am questioning is whether a plausible explanation of the relation between these varied phenomena and indexical interpretation can be had in terms of a singular convention. The problem is that these cases are extremely diverse, and no intention-free description of them seems available to unify them into a single sort of interpretive problem — in response to which a convention might in fact arise.¹¹

One might, however, worry that it seems implausible to think that there was ever a point at which there were multiple ways that the indexicals could be used on devices like answering machines. This, I take it, speaks more to the limitations of our imagination rather than to any fact of the matter. Consider a newer sort of recording: video wills. Both (5) and (6) seem to be perfectly acceptable ways to begin a video will.

(5) Today I met with my lawyer to go over all the details before making this videotape.

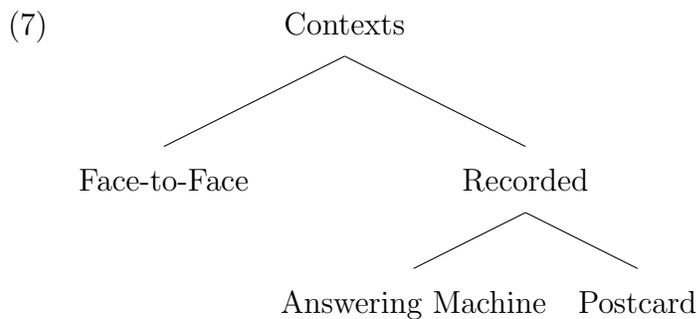
(6) Today you all received a phone call telling you to come to my lawyer’s office.¹²

¹¹This lack of a well-defined set of cues that indicate when ‘now’ refers to a past time looks particularly problematic for Corazza et al. (2002), who do in fact propose that the (intention-insensitive) conventions governing ‘now’ specify that it sometimes refers to times in the past. The worry is that, except via an appeal to intentions, there seems to be no promising response to the question: when exactly?

¹²Both of these examples are from Sherman (2013, p. 21). Sherman takes (5) and (6) to illustrate a phenomenon he calls ‘indexical choice’ — one interesting upshot of which is that, subsequent to an initial choice-point, the speaker must remain faithful to the ‘frame’, or way of conceiving of the objects in the context, that she has adopted. For Sherman, the semantics of indexicals is best thought of in terms of such framing effects, not in terms of rules of reference. While I am generally sympathetic to Sherman’s approach, I worry that his theory leaves unexplained why the indexicals are, by default, interpreted in only certain

Yet (5) and (6) use ‘today’ (another paradigm indexical) to communicate something about two different days — either the day on which the video was recorded or the day on which it is played back. I take it that what (5) and (6) demonstrate is that a singular convention for using ‘today’ on video wills has yet to emerge — if it ever will. This leaves two natural ways of extending the use of ‘today’ to video wills — two different ways of conceiving of the date of the context — open as live-options for the speaker. Should one of these uses of ‘today’ ever become the dominant use on video wills, the character-shifting theory offers a way of formally characterizing the change that natural language would have just witnessed. We would introduce a new type of context and an accompanying character-rule.

Before moving on, it will help to offer a visual map of the picture I am proposing. Kaplan’s theory, the reader will recall, recognizes only one sort of context (contexts of utterance) and one sort of character-rule (rules for that sort of context). I propose, on the other hand, that we relativize character-rules to different sorts of contexts. I offer the following provisional typology of contexts:



This is to be coupled with the following guide to the character-rules for each of these types of contexts:

ways on many sorts of recordings — like answering machines and postcards — even though alternative interpretations might also be possible. There may, however, be ways of combining our theories so as to account for both of these effects.

- (8) Face-to-Face ‘I’ refers to the speaker
 ‘here’ refers to the location of production
 ‘now’ refers to the time of production
- Answering Machine ‘I’ refers to the owner of the line
 ‘here’ refers to the location of playback
 ‘now’ refers to the time of playback
- Postcard ‘I’ refers to the author
 ‘here’ refers to the location of production
 ‘now’ refers to the time of production

In short, according to my character-shifting theory, indexicals are to be evaluated at a context of interpretation according to a rule that is itself determined by what sort of communicative channel is used to deliver the message at that context (the feature by which I type contexts of interpretation). One set of rules governs the reference of indexicals in face-to-face contexts, whereas a different set of rules determines reference when indexicals are played back on answering machines or written on postcards.

4 Three Objections

I now turn to three sorts of concerns facing the character-shifting theory of indexicals. Working through each of these will, I hope, serve to clarify the view.

4.1 Ad Hockery

On first encountering my character-shifting theory, one might worry that it is viciously ad hoc. Am I not just proposing that we shift the meanings of indexicals willy-nilly so as to make them match our untutored intuitions on what those terms refer to in particular contexts?

No. In empirically-driven philosophy of language, we are ultimately responsible to our best data — and, traditionally, competent language users’ intuitions on the truth and falsity of well-formed tokens of natural language have been thought to be amongst those data. But I am not relying here on just these intuitions. I am relying on *both* these intuitions and on observed regularities in the communicative effects of those tokens (that is, what interpreters standardly take away from recorded tokens of (1)–(4)). If responsiveness to such data is sufficient to make a theory ad hoc, then my proposed theory is certainly ad hoc. I would, however, vigorously dispute the claim that this ad hocery is anything but virtuous.

But there is more to be said on this topic. First, the fact that a semantic theory predicts the existence of a number of regularities that require further empirical investigation in order to be fully characterized should be viewed as a weakness only if there is reason to doubt that these regularities actually obtain. Otherwise, the theory should be viewed as grounding a new research program. I contend that my character-shifting theory of indexicals promises to be the start of just such a program.

Second, my character-shifting theory follows quite naturally from a plausible story regarding the development of the meaning of indexicals in natural language. According to the character-shifting theory, the rules governing face-to-face contexts (which are very much like Kaplan’s classical characters) are to be considered the *basic* rules of reference corresponding to each of the indexicals. This claim derives from a view on which face-to-face, voiced language-use is considered more basic than recorded uses of language. This is both a historical and a psychological claim: humans began to voice language long before we began to record it, and most of us learns to speak before we learn to read, write, or otherwise deal with recorded language. The development of conventions that shift the character of an indexical in certain types of contexts thus mark a major achievement in the development of human language — they indicate a juncture at which human beings collectively learned, and humans individually learn, to adapt their use of certain key terms of their languages to

a new and interestingly different sort of communicative environment.

4.2 Bad Predictions

One might also worry that the character-shifting theory still makes some bad predictions. Consider, for instance, an answering machine on which (9) is recorded.

(9) I am here now, of course, but not when you're hearing this.¹³

According to the character-shifting theory, since this is an answering machine context, (9) should be true iff the owner of the line is present at the time and place of playback, but is not present when the listener is hearing that playback. So (9) is predicted not only to be false, but also contradictory.

While I grant that we are capable of hearing (9) as potentially true, the interpretation of (9) in this manner strikes me as nowhere near as smooth and unreflective as the interpretation of recorded tokens of (1). In fact, it is just this markedness that seems to make (9) humorous. Such markedness also makes (9) a good candidate to be explained in terms of 'accommodation' (cf. Lewis (1979)). The idea is that playback tokens of (9) are literally both false and absurd. However, the availability of a nearby interpretation of the indexicals 'here' and 'now' that serves, potentially, to make (9) true prompts us to prefer this non-standard interpretation to the standard, absurd one. Plausibly, such pragmatic repair is only possible if the desired interpretation involves a more basic character-rule than would be standard in the context-type in which the indexical is tokened. Best I can tell, this posit effectively avoids over-generation worries while allowing us to explain the full range of actual cases.

One might worry, however, that this response cuts against the spirit of the character-shifting view: given that (9) can be interpreted as non-absurd and potentially true, why not fold this interpretation into our semantic theory? I prefer not to do this for two reasons. First,

¹³I am told that John Perry once had an answering machine message very much like this.

it significantly complicates the theory for what is, at best, a minor payoff. Uncontroversially, (9) is a marginal case. Second, one important desideratum of semantic theory is to capture the strong default interpretations of terms like indexicals. If we make indexicals ambiguous in answering machine contexts, we lose our best explanation of how such defaults arise. Again, this isn't to close off the possibility of using 'here' and 'now' in the way exhibited in (9) — it's just to claim that our ability to innovate with language in ways like this isn't the sort of thing that semantic theory aims to capture. Plausibly, semantics aims to account for conventionalized regularities that arise in language use — the sort of thing that should give rise to a strong default interpretations — not all of the vicissitudes of actual uses of indexicals.

4.3 Ambiguity

Finally, one might worry that the character-shifting theory makes indexicals ambiguous — but that, plausibly, they are not. This objection, however, is off the mark. Kaplan's semantics for indexicals allows their reference to vary with context. Yet that theory is not standardly understood as entailing that indexicals are ambiguous, since it offers a unitary characterization of their meaning via the context-invariant notion of character. My character-shifting theory functions analogously: 'I' refers to the agent, 'here' to the place, and 'now' to the time, at all contexts. What is allowed to shift is what it is to be the agent, place or time of a context; these can vary with the type of context we are dealing with.¹⁴ Alternatively, we can think of the present theory as associating indexicals not with simple characters, but rather with a constant 'meta-character': that is, a function from a context-type to an ordinary character, in the broad sense (e.g. 'I refers to the speaker/author/owner of the line' as opposed to just 'I refers to the agent'). While these ordinary characters will shift relative

¹⁴Note that Kaplan was already doing something similar to this when he claimed that what it is to be the agent is to be *either* the speaker *or* the author (1989b, p. 505).

to context-types, this meta-character remains constant.

Crucially, on both these ways of understanding the present proposal, multiple uses of an indexical at a single context will always refer to the same thing. The problematic cases that have been pointed out in the literature, and which have been taken to tell against the indexicals being ambiguous, all depend on shifting reference within a single context (and, hence, context-type) — something which the character-shifting theory prohibits.¹⁵

Typing contexts also allows the character-shifting view to preserve Kaplan’s truisms about face-to-face utterances. Kaplan made (3) a logical truth and (1) a logical falsehood in order to mark the special relation that obtains between utterances of these sentence-forms and the world in cases of spoken language. Once we consider recorded tokens of these sentence-forms, however, such generalizations fail — at least so long as we take our truth-judgments at face value. The character-shifting theory, however, allows us to retain something close to the original notion of logical truth while still accounting for answering machine tokens of (1) by introducing the concept of a truth *relative to a context-type*. So, (3) is true, and (1) false, relative to face-to-face contexts — since (3) is true and (1) false at all such contexts. Similar to Kaplan’s original theory, this claim is supposed to hold in virtue of the special relation between agents, times, and locations in face-to-face utterance contexts. This special relation fails to obtain, however, with respect to other context-types — making it incorrect to characterize these sentences as *logically* true or false.

¹⁵ Cohen (2013) aptly demonstrates why positing ambiguity within a single utterance ought to be avoided in his discussion of Smith (1989), who posits that indexicals are straightforwardly ambiguous. Smith’s theory predicts, for instance, that there should be a true reading, in a single context and without Alyosius moving, of “Alyosius is here but Alyosius is not here” — namely, when this message is recorded in Alyosius’ absence but played back in his presence (this example is from Cohen (2013, pp. 5–6)). I concur with Cohen that a true reading is impossible here. My own theory avoids positing a true reading here by fixing a single character-rule for ‘here’ relative to a context-type.

5 Cohen and the Context of Tokening

Cohen (2013) offers a minimal modification of Kaplan’s semantics aimed both at dealing with the answering machine problem and with accounting for Kaplan’s observations regarding (1) and (3). Specifically, Cohen claims that ‘here’ and ‘now’ are sensitive to the context of *tokening*, whereas ‘I’ is sensitive to the context in which it was *originally uttered or inscribed* (Cohen 2013, p. 10). So, in the answering machine case, ‘I’ refers to the speaker who recorded the message (since she made the original utterance), whereas ‘here’ and ‘now’ refer to the place and time of playback, respectively. In face-to-face communication, these contexts will be identical. Thus, so long as we restrict ourselves to considering only face-to-face utterances, there is no harm in collapsing these two types of contexts — in effect re-deriving Kaplan’s semantics for indexicals for this type of context (Cohen 2012, pp. 10–11). This collapse allows Cohen to preserve Kaplan’s truisms regarding face-to-face contexts.

Cohen’s theory handles the answering machine case with ease — predicting that each time the answering machine is played back and (3) is tokened, that recording-token is true iff the original speaker is not at the time and place where the machine is triggered.¹⁶ Unfortunately, Cohen’s theory stumbles on the postcard case. This is because the theory predicts that ‘here’ and ‘now’, when written on postcards, refer to the place and time where the card is ‘tokened’. What exactly constitutes a tokening in the postcard case is a bit less clear than in

¹⁶Romdenh-Romluc (2006, pp. 262–63) offers a variant on this case where these truth-conditions plausibly fail to match our intuitions on what is communicated by the message. The case hinges on the owner of the line having asked someone else to record her message for her. Such ‘friend cases’ are interesting, but I am not convinced they should be dealt with semantically. I therefore leave this objection to the side for the purpose of evaluating Cohen’s proposal. Two things about such cases are worth briefly noting: first, my particular variant of character-shifting theory can deal with such friend cases semantically. However, this doesn’t strike me as a central consideration in favor of the theory. Second, on Cohen’s theory, Romdenh-Romluc (2006)’s friend case actually turns out to be a Gettier case — since the listener will come to believe that the owner of the line isn’t at home on the basis of coming to believe that the recording of (1) is true, and (1) is in fact true iff the owner’s friend isn’t at the owner’s home at the time of playback. Assuming that neither the owner of the line nor the owner’s friend are at the owner’s home when the answering machine is triggered, the listener will have justifiably formed a true belief on the basis of antecedently accepting a true, but irrelevant, content. In spite of such justification, I take it that the listener does not know that the owner is not at home.

the answering machine case — since we are no longer dealing with a device that is literally triggered and played back. The most straightforward analogue, it strikes me, would be to treat each instance of a postcard’s being read and interpreted as constituting a tokening.¹⁷ If that is correct, then Cohen’s theory predicts that each token of (4) is true iff it is beautiful at the time and place where the postcard is read. But that hardly seems like the right prediction.

Cohen (2013) recognizes the threat posed by such cases and proposes a clever response: he claims that the terms ‘here’ and ‘now’, as inscribed on postcards, aren’t actually indexicals. Rather, they are either bound variables — with reference determined by the place indicated on the postcard and the time written or stamped on it, respectively — or they are (unbound) demonstratives (Cohen 2013, p. 24 fn. 32).

Either version of this response faces substantial difficulties. First, let us consider the suggestion that postcard uses of ‘here’ and ‘now’ are actually demonstrative uses of those terms. I accept both (i) that there are demonstrative uses of ‘here’ and ‘now’ and (ii) that we have yet to discover a clear linguistic test for distinguishing indexical from non-indexical uses of ‘here’ and ‘now’. Thus, I do not want to object to Cohen on the grounds that this move is ad hoc. Instead, I object to his suggestion because, if these are demonstrative uses of ‘here’ and ‘now’, they should be sensitive to the speaker’s intentions in a way that these particular tokens do not appear to be. It just doesn’t seem to matter, for instance, what Tarek intended when he wrote ‘here’ on the postcard — that token of ‘here’ refers to the place where he inscribed it.¹⁸

Second, let us consider a bound variable account of postcard uses of ‘here’ and ‘now’.

¹⁷In defending his view against an earlier version of this objection, Cohen seems to concur (Cohen 2013, p. 24 fn. 32).

¹⁸Cohen might claim at this point that our ordinary intuitions actually track pragmatic, rather than semantic, content in this case. That is, Tarek’s use of ‘here’ *really* refers to whatever he intends it to refer to; we just tend to confuse ourselves about this. However, such a move would be inconsistent with Cohen’s own (convincing, I think) objections to Predelli (2002)’s invocation of the semantics/pragmatics distinction in an analogous context, in response to a very similar set of challenges (Cohen 2013, pp. 11–12).

My objection to this suggestion is that it still makes incorrect predictions. Imagine once more that Tarek has mailed me a postcard with *just* (4) inscribed on it. Imagine further that it was beautiful at the place/time where he inscribed the card. However, suppose that he didn't actually get around to mailing the card until several weeks later — by which point the weather had taken a serious turn for the worse. The postcard is postmarked at that later date, and that postmark provides the only plausible binder for the variable associated with 'now'. But then Cohen's theory predicts that (4), as inscribed on this postcard, is false — which seems like the wrong result. What's more, this version of Cohen's theory seems to predict that the truth-conditions of a completely undated, unstamped postcard should be difficult, if not impossible, to elaborate — since nothing would be available to bind the relevant variables. But, while undated, unstamped postcards certainly present interpretive difficulties, their truth-conditions are not, in fact, hard to describe. 'I' just refers to the author, 'here' to the place of inscription, and 'now' to the time of inscription.¹⁹

In short, Cohen's theory is admirable in its attempt to preserve both Kaplan's truisms and much of the simplicity of Kaplan's original semantics. However, it is far less clear that Cohen's theory can deliver empirical results that match the character-shifting theory in terms of their accuracy. Attempts to reconcile the theory with the data look bound to introduce significant additional complexity — thereby undercutting the simplicity of the theory that is one of its main virtues. I suggest, therefore, that we ought to prefer my character-shifting theory of indexicals over Cohen's context of tokening approach.

¹⁹Another option for preserving Cohen's view, undiscussed by Cohen himself, would be to claim that indexicals are actually ambiguous. However, Cohen explicitly denies this (Cohen 2012, pp. 8–9). What's more, he advertises the fact that he can preserve indexical-univocality as one of the primary virtues of his theory.

6 Predelli and Context-Shifting

Predelli (1996, 1998a, 1998b, 2002, 2005) offers a semantic theory capable of accounting for both the answering machine and postcard cases, but which explicitly gives up on Kaplan's truisms regarding face-to-face utterances. Basically, Predelli advocates a position on which speakers can intentionally shift the context relative to which the indexicals they use ought to be semantically evaluated. In other words, pace Kaplan, Predelli claims that indexical reference is intention-*sensitive* — not in virtue of the character-rules themselves being intention-sensitive, but in virtue of speakers' intentions determining, relative to any use of an indexical, which person, place, and time serves as the agent, place, and time of the context (Predelli 1996, pp. 88–89).

Here is how Predelli's theory handles answering machine cases: when speakers record outgoing messages, they typically intend for their uses of 'here' and 'now' to be evaluated relative to each context in which the machine is triggered; 'I', on the other hand, is standardly intended to be evaluated relative to the original context of utterance. These original intentions endure over the lifetime of the recording, determining reference at each instance where the machine is triggered. In the postcard case, on the other hand, writers standardly intend for both 'here' and 'now' to be evaluated relative to the context of inscription, not the context of evaluation. Such presumed intentions explain the content standardly communicated in both sorts of cases; this content is smoothly communicated because we all know which intentions standardly accompany uses of 'I', 'here', and 'now' on both answering machines and postcards. Importantly however, it is the speaker's *actual* intentions in such cases that determine the reference of the indexicals in — and, ultimately, the truth-values of — recorded tokens of (1) and (4). So, according to Predelli, semantic content can detach, quite radically even, from the content ordinarily communicated by a token of an indexical-containing sentence-type.

There are several reasons to be wary of Predelli's proposal: first, by embracing improper contexts (which Predelli does explicitly at Predelli (2002, p. 311)) and rejecting the claim that any regularities obtain in the reference of indexicals, even relative to particular context-types, Predelli gives up on both logical truth and on weaker, related notions — like my truth relative to a context-type. Predelli is likely to call this a feature of his view rather than a drawback, so I do not mean to put too much weight on this objection. Still, both Cohen and I have demonstrated that we do not in fact face a simple choice between preserving something like logical truth and dealing with the answering machine problem. It ought, therefore, be incumbent on Predelli to demonstrate why we should give up on the idea that (1) and (3) exhibit a special sort of relation to the world — at least relative to certain types of contexts. Second, on reflection, it is not at all clear that we need to know anything about the speaker's intentions in order to understand the meaning of a recorded token of (1). Third, Predelli's theory over-generates predictions of successful context-shifting.

These latter two objections are in fact closely related. To see how, consider a concrete case: suppose that my office-mate is out sick today and her students are massed outside our door waiting for what would ordinarily be her office hours. I open the door and shout (1) while intending to shift the context of evaluation to one in which my office-mate, not me, is the agent. What is clear about this situation is that my utterance of (1) is highly unlikely to succeed in communicating to these students that my office-mate is not here today. What is less clear is whether this observation is semantically significant.

Predelli's theory claims that, so long as I had the right intention here, I succeeded in shifting the context relative to which (1) ought to be evaluated. So, my utterance of (1) is true iff my office-mate isn't here today. *Prima facie*, this seems like the wrong thing to say about the utterance.²⁰ Predelli can, of course, contest the significance of this intuition by appealing to the distinction between semantic and pragmatic content. In fact, he does

²⁰The basic form of this objection is derived from Corazza et al. (2002, p. 9).

just that at Predelli (2002, pp. 314–15): according to Predelli, when indexicals are involved, communicated content is a poor guide to semantic content. This is because what is communicated tends to depend on listeners' *assumptions* regarding the speaker's intentions rather than what the speaker actually intended. Since we are so used to trying to figure out what speakers must reasonably be intending when they utter sentences at contexts, we can easily be misled into thinking that utterances like (1) have truth-conditions that are independent of the speaker's actual intentions. This, according to Predelli, is a natural mistake that can only be corrected by extreme vigilance on the part of the semantic theorist.

The problem with this response is that it serves to undercut the very evidence that would lead us to prefer Predelli's theory of indexicals over Kaplan's original semantics in the first place. Consider that, by an analogous set of considerations, we could just as well hew to an unmodified Kaplanian semantics for indexicals in the face of the answering machine problem. We would begin by noting once more that communicated content is a poor guide to semantic content — at least for cases involving indexicals. Then we would claim that the content standardly communicated by (1) when played back on an answering machine can be explained by appealing to what listeners tend to assume about how speakers intend to use sentences like (1) in such contexts. But, according to this line of reasoning, this tells us nothing directly about semantic content. Given that Kaplan's semantics is undoubtedly simpler than Predelli's alternative, we would seem to have little or no reason at this point to prefer Predelli's theory over Kaplan's. In other words, in order to motivate a departure from Kaplan's original semantics, Predelli had to claim that semantic content should generally match communicated content for tokens of (1) when played back on answering machines. Once we accept that communicated content is a poor guide to semantic content, we have no reason to think that Kaplan's semantics needs to be revised in light of answering machine cases. Predelli's theory, amended so as to deal with the over-generation worry that it faces,

thus turns out to be self-undermining.²¹

7 Recording Pragmatics

The failure of Predelli’s theory thus raises the possibility of pursuing an entirely different sort of option: accounting for our truth-judgments regarding (1) by appealing to pragmatics rather than by modifying Kaplan’s original semantic theory. I will argue here that (i) the most obvious pragmatic option faces a serious technical challenge, and (ii) there are some more general reasons for thinking that what’s needed to account for answering machine and postcard cases is actually a semantic, rather than a pragmatic, explanation.

First, the technical challenge: I assume that the most obvious way to generate a pragmatic account of answering machine recordings of (1) is to claim that, while their semantics is properly given by an unmodified Kaplanian account, they *con conversationally implicate* some other content — and that that implicated content will often be true (cf. Grice (1989b)). That is, while answering machine tokens of (1) 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.²²

The problem with this suggestion is that, at least according to Grice, conversational implicatures should be ‘cancelable’ (Grice 1989a, p. 44) — and this purported implicature is not. Consider the following pair:

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

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

An utterance of just the first clause of (10), ‘Laura and Danny had a kid and got married’,

²¹Cohen (2013, pp. 13–15) levels a related objection against Predelli’s appeal to the semantics/pragmatics distinction.

²²This pragmatic strategy thus consciously mirrors the structure of Kripke (1977)’s response to Donnellan (1966)’s claims regarding the semantics of definite descriptions.

standardly communicates that they did so kid-first. However, there is nothing contradictory about an utterance of (10) — 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 the semantic content expressed by an utterance of (10). On the other hand, when (11) 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 (1) 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 (11) should be non-contradictory — since the first and second clause would be about different times, and possibly about different places as well.²³

Now for the more general theoretical worry: pragmatic solutions 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. When (1) is played back on an answering machine, understanding that the ‘here’ and ‘now’ refer to the place and time of playback is entirely non-optional. If one fails to interpret these indexicals in this way, one hasn’t just failed to ‘get it’, nor has one been a bit slow on the uptake. Rather, one has *failed to understand* the sentence-token at a fairly basic level. Linguistic conventions look to be the right sort of things to ground strong demands on interpretation like this one. A listener is responsible for knowing the conventions of her linguistic community, and a failure to grasp these can lead to a basic failure to understand an

²³In light of these issues, one might be tempted to instead classify the content conveyed by (1) as recorded on an answering machine as conventionally, rather than conversationally, implicated. This suggestion faces two problems: first, it is not at all clear that conventional implicature is actually a species of pragmatic content. Grice calls it a type of ‘conventional meaning’, seemingly in contrast to conversational implicature (Grice 1989c, p. 121). Since Grice never uses the term ‘semantic content’, it is difficult to say how exactly we ought to map that notion onto his various sorts of meanings. Semantics, however, is often conceived of as tied to conventions as opposed to use — making conventional meaning more likely a species of semantic, as opposed to pragmatic, content. Second, conventional implicatures generally ‘add to’ the total content communicated by an utterance, rather than supplanting the semantic content associated with that utterance (cf. Neale (2001)). Here, however, it seems odd to think that answering machines convey *both* something absurd about the speaker’s state at the time and place where she made the recording and something potentially true about the speaker’s state at the time and place where the answering machine is triggered.

utterance or recorded token of language — not just a failure to understand all the exigencies of how that utterance or recording token was being used in some particular context.

Put slightly differently, pragmatics would seem to be in the business of characterizing how we can leverage conventional meaning in order to put utterances and recordings to various uses in particular circumstances. But such characterizations, while certainly interesting, aren't robust enough to ground norms regarding the basic interpretive standards that a listener must live up to in order to qualify as a competent language-user. Rather, to ground such basic interpretive norms — norms that determine what is required to achieve basic understanding of a sentence-token, relative to a context — we need to appeal to conventions governing the ordinary use of particular terms and syntactic forms. This places us squarely in the realm of semantics. Now, given that a failure to understand 'here' and 'now' as recorded on an answering machine as referring to the place and time of playback, respectively, would seem to constitute a full-blown interpretive failure, we have good reason to posit that there are interpretive norms at work here — grounded in conventions that assign each of these indexicals to its respective referent, relative to this type of context. But, if that's right, then these rules are properly conceived of as semantic.²⁴

Pragmatic accounts of recorded indexicals, therefore, won't be nearly as simple as initially advertised. First off, such accounts will have to specify how the content standardly communicated by answering machine tokens of (1) should be classified — that is, as just which sort of pragmatic content. This looks to be no easy task. In addition, the pragmatic theorist will need to explain both (i) what mechanisms serve to convey that content, and

²⁴This is essentially a variant of Devitt (1997a, 1997b, 2004) and Reimer (1998a, 1998b)'s 'Argument from Convention' against unitary Russellian theories of definite descriptions: pragmatic explanations are well-suited for explaining why we might interpret such-and-such word or phrase in a particular way on one or another particular occasion. They are, on the other hand, significantly less well-suited for explaining why people regularly and consistently interpret a word or phrase in a particular sort of way across a wide range of contexts — and why they are, in fact, justified in doing so. Even if such an interpretive regularity were not originally the result of a convention, it would seem that a convention would be expected to arise over time, supposing that this interpretive regularity were sustained for a sufficient period.

(ii) how those mechanisms give rise to judgments that listeners exhibit a particularly serious sort of incompetence when they fail to interpret ‘here’ and ‘now’, as tokened on answering machines, as referring to the place and time of playback. In contrast, the character-shifting account of indexicals can appeal to the conventions associated with those indexicals, relative to answering machine contexts, to explain these strong default interpretations. Competence with English, extended to include the recording and interpretation of answering machine messages, is partly constituted by interpreting ‘I’ as the owner of the line, ‘here’ as the location of playback, and ‘now’ as the time of playback when those indexicals are tokened on answering machines.

8 Conclusion

In summary, I have argued that the answering machine problem is best dealt with via a semantic approach. In particular, I have argued in favor of a character-shifting theory of indexical meaning, both as against other semantic accounts and as against the pragmatic strategy. There is, however, one last defense of Kaplan’s original semantics for indexicals that I have not yet considered: perhaps Kaplan’s theory is meant to account only for utterances, and not for well-formed tokens of natural language more generally. The theory would then bear no responsibility for accounting for the meanings of tokens of natural language as recorded on answering machines and postcards — since, plausibly, these aren’t utterances. Just such a defense of Kaplan’s semantics has in fact been offered in Stevens (2009).

Exegetically, this response is implausible. In “Demonstratives” (1989b), Kaplan speaks primarily of evaluating ‘sentences in contexts’, not of evaluating ‘utterances’. Only, later, in “Afterthoughts” (1989a), does Kaplan begin to speak primarily in terms of utterances. And, whether or not they qualify as utterances, sentences as tokened on answering machines and postcards are clearly sentences in contexts.

Regardless of its exegetical accuracy, we might still ask whether such a defense of an utterance-version of Kaplan’s semantics is persuasive. I contend that it is not. While Stevens’ suggested defense is certainly coherent, there is little to no philosophical payoff to be had from it. We are still left facing the answering machine problem. It just turns out to be a problem not for Stevens’ ‘semantics’ — which is defined only over utterances — but rather for ‘semantics*’ — the study of the conventional meaning and truth-conditions of well-formed tokens of natural language on whatever sort of medium they happen to appear. Stevens’ semantics would indeed constitute a sub-part of semantics* in which the answering machine problem does not arise. But this hardly resolves the problem — particularly since semantics* promises to be the richer and more interesting of these two projects.²⁵

In spite of its problems, Stevens’ suggested defense does serve to highlight what I take to be the basic methodological upshot of this paper: in initially generating semantic theories, it will often be helpful to constrain the data under consideration. One way to do this when considering indexicals is to ignore recordings and focus solely on face-to-face contexts, or utterances. We must take care, however, not to let these initial restrictions take too deep a root. We may, at times, be tempted to preserve the simplicity and elegance of our initial theories by pushing data that cuts against them into the ‘waste bin of pragmatics’. Sometimes, this will be appropriate, but often it will not. Once we note that, at this point in human history, recorded uses of language clearly aren’t aberrant, the burden falls on the theorist who wants to treat recorded tokens as somehow ‘less semantic’ than voiced tokens of natural language to explain why such a differential treatment is justified. Recorded

²⁵There is an additional worry about Stevens’ proposal as well: the proposal relies on a sharp distinction between saying and asserting according to which, when (1) is played back on an answering machine, nothing is asserted (Stevens 2009, pp. 217–19). Kaplan’s theory, by Stevens’ lights, provides an account of asserted content, something which is only properly associated with utterances, not recordings. Presumably though, something is still said by (1) when tokened on an answering machine. Unfortunately, Stevens fails to specify what exactly he takes this content to be. Even more problematically, the rough gloss he does provide looks to be at odds with Kaplan’s original semantics for indexicals (Stevens 2009, p. 217). This leaves one to wonder both just what Stevens’ theory of the said content of indexicals is, as well as how that said content is supposed to help explain the content asserted by uses of sentences containing indexicals.

tokens of natural language exhibit the same sorts of regularities on which semantic theories have standardly relied — including strong, normatively significant default preferences for the basic interpretation of tokens of a recognizable morphological type. We therefore have strong *prima facie* reasons for treating observed regularities in recorded uses of the indexicals as semantically significant.

There is, however, a nearby claim that we would do well to endorse: while widespread, recorded uses of language are also *non-basic*. Spoken language both long predates recorded language and is what each of us typically learns first. Clearly, it can also exist in the absence of a written counterpart. Spoken language, therefore, has a good claim to being self-standing: we should be able to give a full, accurate semantics of a spoken language without ever mentioning recorded tokens of that language. The character-shifting theory attempts to respect this basic-ness of spoken language by treating it as a well-defined sub-part of our actual language — a sub-part on which more complicated uses, like recorded uses, are built. The conventions that arise out the use of recording technologies can shift the character of indexicals *from their original character*, as defined by the role they play in spoken language. However, in natural language — at least of the sort that humans have so far developed — shifted-characters cannot exist in the absence of these more basic character-rules.

Recorded language, on this view, is thus quite literally an extension of spoken language. Spoken language is extended by repeated coordination on one of the natural ways that particular indexicals might be used to refer when their ordinary referents are unavailable, or perhaps uninteresting, at the relevant context of interpretation. Such character-shifting is, I believe, just a particular instance of a more general phenomenon — namely, humans' ability to adapt our language to circumstances very different from those in which it arose. Similar shifts in meaning over time, and in response to changed circumstances, can be observed in a very different philosophical context as well: in the 'wandering significance' of numerous mathematical and scientific terms (cf. Wilson (2006)). Endorsing the character-shifting

view of indexical meaning thus amounts to endorsing the claim that, in recorded contexts, indexicals have a tendency to wander. Mapping the full extent of their various journeys will, I hope, prove to be an interesting empirical project.

Acknowledgments

Thanks first to my advisers on this project: David Kaplan and Sam Cumming. Particular thanks are also due to Jonathan Cohen, whose work on indexicals has significantly informed my views. For extensive discussion and comments, thanks to: Zed Adams, Josh Armstrong, Matthew Babb, Kent Bach, Mel Bervoets, Alexis Burgess, Arudra Burra, Alejandro Pérez Carballo, Jorah Dannenberg, Shamik Dasgupta, Kenny Easwaran, Andy Egan, Ashley Feinsinger, Marina Folescu, David Friedell, Nathaniel Hansen, Daniel Harris, Brent Kious, Karen Lewis, Kathryn Lindeman, Dustin Locke, Chauncey Maher, Alex Morgan, Matthew Moss, Sarah Murray, Tyke Nunez, Jessica Pepp, David Plunkett, Alex Radulescu, Andrew Reisner, Indrek Reiland, Brett Sherman, Will Small, and Will Starr. Thanks as well to my audiences at the UCLA Albritton Society, the Rutgers philosophy of language working group, the New York Philosophy of Language Workshop, and the CLAP.

References

- Bach, Kent. 1992. Intentions and demonstrations. *Analysis* 52. 140–146.
- Cohen, Jonathan. 2012. Indexicality and the puzzle of the answering machine. Manuscript (forthcoming in *The Journal of Philosophy*).
- Corazza, Eros, William Fish & Jonathan Gorvett. 2002. Who is I? *Philosophical Studies* 107. 1–21.
- Devitt, Michael. 1997a. Meaning and psychology: A response to Richard. *Noûs* 31. 115–131.
- Devitt, Michael. 1997b. Responses to the Maribor papers. In D. Jutronic (ed.), *The Maribor papers in naturalized semantics*, 353–411. Pedagoska Fakulteta.

- Devitt, Michael. 2004. The case for referential descriptions. In Marga Reimer & Anne Bezuidenhout (eds.), *Descriptions and beyond*, 280–305. Oxford University Press.
- Donnellan, Keith. 1966. Reference and definite descriptions. *Philosophical Review* 75. 281–304.
- Grice, H. Paul. 1989a. Further notes on logic and conversation. In *Studies in the way of words*, 41–57. Cambridge: Harvard University Press.
- Grice, H. Paul. 1989b. Logic and conversation. In *Studies in the way of words*, 22–40. Cambridge: Harvard University Press.
- Grice, H. Paul. 1989c. Utterer’s meaning, sentence-meaning, and word-meaning. In *Studies in the way of words*, 117–143. Cambridge: Harvard University Press.
- Kaplan, David. 1989a. Afterthoughts. In Joseph Almog, John Perry & Howard Wettstein (eds.), *Themes from Kaplan*, 565–614. Oxford: Oxford University Press.
- Kaplan, David. 1989b. Demonstratives. In Joseph Almog, John Perry & Howard Wettstein (eds.), *Themes from Kaplan*, 481–563. Oxford: Oxford University Press.
- Kripke, Saul. 1977. Speaker’s reference and semantic reference. In Peter A. French, Thomas Uehling, Jr. & Howard K. Wettstein (eds.), *Midwest studies in philosophy, vol. 2*, 255–276. Minneapolis: University of Minnesota Press.
- Lewis, David. 1969. *Convention*. Cambridge: Harvard University Press.
- Lewis, David K. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8(3). 339–359.
- Neale, Stephen. 2001. Implicature and colouring. In G. Consenza (ed.), *Paul grice’s heritage*, 135–180. Rome: Brepols.
- Perry, John. 2006. Using indexicals. In Michael Devitt & Richard Hanley (eds.), *Blackwell guide to the philosophy of language*, chap. 17, 314–334. Oxford: Blackwell.
- Perry, John. 2009. *Reference and reflexivity*. Stanford: Center for the Study of Language and Information 2nd edn.
- Predelli, Stefano. 1996. Never put off until tomorrow what you can do today. *Analysis* 56(2). 85–91.
- Predelli, Stefano. 1998a. ‘I am not here now’. *Analysis* 58(2). 107–115.
- Predelli, Stefano. 1998b. Utterance, interpretation and the logic of indexicals. *Mind and Language* 13(3). 400–414.
- Predelli, Stefano. 2002. Intentions, indexicals, and communication. *Analysis* 62(4). 310–316.

- Predelli, Stefano. 2005. *Contexts*. Oxford: Oxford University Press.
- Quine, Willard V.O. 1968. Ontological relativity: The Dewey lectures 1969. *Journal of Philosophy* 65(7). 185–212.
- Reimer, Marga. 1998a. Donnellan’s distinction/Kripke’s test. *Analysis* 58(2). 89–100.
- Reimer, Marga. 1998b. Quantification and context. *Linguistics and Philosophy* 21. 95–115.
- Romdenh-Romluc, Komarine. 2006. ‘I’. *Philosophical Studies* 128. 257–283.
- Sherman, Brett. 2012. Constructing Contexts. Manuscript, University of Rochester.
- Sidelle, Alan. 1991. The answering machine paradox. *Canadian Journal of Philosophy* 21(4). 525–539.
- Siegel, Susanna. 2002. The role of perception in demonstrative reference. *Philosopher’s Imprint* 2(1). 1–21.
- Smith, Quentin. 1989. The multiple uses of indexicals. *Synthese* 78(2). 167–191.
- Stevens, Graham. 2009. Utterance at a distance. *Philosophical Studies* 143. 213–221.
- Vision, Gerald. 1985. ‘I am here now’. *Analysis* 45(4). 198–199.
- Wilson, Mark. 2006. *Wandering significance: An essay on conceptual behavior*. Oxford: Oxford University Press.