Chapter 17

Using Indexicals

John Perry

1 Introduction

In this essay I examine how we use indexicals. The key function of indexicals, I claim, is to help the audience find supplementary, utterance-independent, channels of information about the object to which or to whom the speaker refers.

This exploration of the use of indexicals is based on the reflexive-referential theory of the meaning and content of indexicals and other referring expressions (Perry 2001).

I review the reflexive and referential aspects of indexicals in sections 2 and 3. In section 4 I explore the way we use indexicals and then, in section 5 try to use the lessons learned to think about some problems cases, suggested by Stephano Predelli and Varol Akman, that suggest that the standard semantic rules for indexicals may be too simple.

2 Indexicals

Icon, index, and symbol

The term 'indexical' comes into the philosophy of language from Charles Sanders Peirce. Here is an explanation of Peirce's threefold division of signs: icon, index, and symbol:

> Signs are icons, indices (also called "semes"), or symbols . . . accordingly as they derive their significance from resemblance to their objects, a real relation (for example, of causation) with their objects, or are connected only by convention to their objects, respectively. (Burch 2001)

Suppose I am talking to Mr. Fritchey. I may use his first name, and refer to him as 'Elwood.'[1] Here, at first pass, it seems that the story goes as follows. There is a completely arbitrary convention that allows me to refer to a certain individual, standing before me, as 'Elwood.' Although in this case he is standing before me—a "real" relation—that has nothing to do with the fact that I refer to *him* when I say "Elwood."[2] I will be able to refer him by using 'Elwood' after he has gone. The work of securing reference is done by the convention, and does not depend on any further connection between the speaker and Elwood. So 'Elwood,' it seems, is a symbol, connected to Mr. Fritchey only by convention.

One of Peirce's central examples of indexicals was smoke, which is a sign of fire. Here the real relation is causation: fire causes smoke. No convention is involved, and no language. Smoke is a natural sign of fire, not a conventional one. The term "indexicality" is not much used for natural signs in contemporary philosophy of language, however.

Since I am talking to Elwood, I can refer to him with the word 'you.' The word 'you' has a conventional meaning in English. We use it to refer to the person to whom we are talking. However, this conventional meaning does not determine reference all by itself. 'You' refers to Elwood, *given* its conventional meaning, because Elwood is *the person I am talking to*. This is a real relation between Elwood and I, involving causation and perception. This conventional species of indexicality is our topic, and it is the phenomenon for which "indexicality" is used in contemporary philosophy of language. Words like 'here,' 'now,' 'you' and 'I,' with their partly conventional, partly relational, links to their referents, are paradigms.

Utterances and tokens

The term 'token' is also due to Peirce. Tokens are distinguished from types; tokens are particular bursts of sound or bits of ink. Consider the list: cow, dog, cow. It contains *two* tokens of the type *cow*, one token of the type *dog*.

Hans Reichenbach developed his token-reflexive account of indexicals in *Elements of Symbolic Logic* (Reichenbach 1947). The expressions in the formal languages studied by symbolic logic up to that time were taken to be non-ambiguous and also had a property Reichenbach called *equisignificance*. This means that two tokens of the same type have the same semantic value. The symbol '2' always stands for the number two; the description 'The first president of the United States' always stands for George Washington. If the expression is a declarative sentence, all tokens of it have the same truth-value. If '$2 + 2 = 4$' or 'All ravens are black' are true when I say them, they will be true when you say them, as long as we are using the same words with the same meaning.

In contrast, indexicals and larger expressions containing them are not equisignificant. Different tokens of the same type, with the same meaning, can stand for different things. The type gives us a relation between tokens and semantic values.

Hence, different tokens of the same type can have different semantic values. Different tokens of 'I' stand for different people; different tokens of 'here' stand for different places, and so forth. Elwood might say truly

(1) I have been to Paris

while Elwood Jr. when he repeats his father's remark,

(2) I have been to Paris

says something false. They both use the same English sentence, with the same meaning, but their statements have different truth conditions and different truth values.

Reichenbach sees that we could get at what (1) and (2) have in common and how they differ, by stating the truth-conditions in terms of the two different tokens:

(3) (1) is true IFF the speaker of (1) has been to Paris.
(4) (2) is true IFF the speaker of (2) has been to Paris.

The type gives us a formula, the same for every instance; the token gives us the thing to which the formula applies. I call this *reflexive* because the token *itself* is mentioned in giving the truth conditions.[3]

I emphasize the distinction between *tokens* and *utterances*. By 'utterance' I mean an intentional act of speaking, signing, typing, writing, etc. By 'token' I mean an effect of such acts, a burst of sound or a mark that is intended to be perceived, recognized, and interpreted by a hearer or reader. In *Elements* Reichenbach confuses or conflates utterances and tokens. He announces that he means an *act* by 'token,' but soon is talking about the *ink marks* on a page as tokens.

Utterances are *semantically* basic. The intentionality of linguistic acts is a special case of the intentionality of purposeful action. The language to which a token belongs, the identity of the words and their meanings, the syntax, the reference of terms, all derive from the minds of the speakers, and connections between those minds, other minds, things, and properties. On the other hand, tokens are often *epistemically* basic. When you read this essay, for example, you see tokens produced by me (not directly by me, of course, but by a complex process I initiated). You take these to be the result of utterances by me—purposeful acts of typing, in this case. When the utterance itself cannot be observed, tokens are what the reader or listener has as evidence. Skilled speakers take into account the extent to which the token will be the main source of information. One speaks louder when the audience is distant, and less loud when they are close. When speaking on the telephone, or writing a note, one should not rely on contextual clues that the hearer or reader cannot perceive.

Tokens can be re-used, and when they are, the new utterances may have different semantic and syntactic properties than the originals. An expensive name-tag that said ‘George Bush’ may have been put in the White House Museum when George I left office; it may now be recycled for George II. A veteran protestor might use a sign, ‘You are a scumbag’ time after time, referring to different politicians. Eros Corazza imagines (or remembers) a philosophy department, faced with steadily diminishing budgets, using a Post-it note, ‘I’ll miss my office hours today,’ for different faculty members on different, days, thriftily getting years of use from the same piece of paper (Corazza 2004)

Wilfred Sellars and others have called utterances “tokenings.” I prefer ‘utterance,’ perhaps because ‘tokenings’ suggests that tokens are semantically basic.

The importance of the utterance/token distinction grows with changes in the technology of language. In face-to-face communication, the token is the burst of sound that travels to one’s ears; there is typically not much difference between perceiving the utterance and perceiving the token. (There is even less in face to face signing, as with American Sign Language.) Writing makes a dramatic difference; tokens remain long after the utterance. Publishing permits the reproduction of tokens; telephony the distant perception of tokens at the time the utterance occurs. Each change in technology makes new patterns of production and perception of tokens possible, and so new expectations and intentions based on these possibilities. We will return to this theme below.

This picture of language use suggests the importance of theories of utterances and tokens. Communicating involves causing physical events that have predictable effects that we can exploit. As noted above, we plan our utterances paying attention to the circumstances under which the tokens we produce will be perceived. An adequate theory of these plans requires representation of the myriad of relations into which utterances and tokens can stand to other concrete objects, people, purposes, projects, and other factors. These factors will figure in the process of reaching a reasonable interpretation of what the speaker is trying to communicate. The most natural way to approach meaning and content would seem to be, then, as properties of utterances.

In David Kaplan’s system (Kaplan 1989) pairs of sentence types and contexts model utterances. Although his theory is in the token-reflexive tradition, it has neither tokens nor utterances. Nevertheless, utterances are implicit in the theory. Kaplan’s contexts model properties of utterances. The quadruple of speaker, time, location, and world are made up of the person, time, place, and world that play the appropriate roles in relation to an utterance.

Kaplan abstracts these properties from the utterance, and combines them with the character of the expression uttered to give us a sentence in context, a pair of context and character. Context and character suffice to determine content, and it is the interplay between context and content on which Kaplan bases his logic. This approach has provided considerable insight about the meaning (character), content, and logic of indexicals and demonstratives. But to understand how indexicals work in communication, why they are useful, and how we develop communication plans using them, we need the utterances and tokens with all of the properties that are involved when one produces them and sees or hears them. The meaning of the sentence used is important, but so may be the volume with which it is spoken, the direction from which it comes, the visibility of the speaker to the hearer, the differences between synonymous expressions, and many other things.

Indexicals and demonstratives

By *indexicals*, then, I mean expressions that have conventional meanings that associate them with certain relations objects may have to utterances of them, or, in somewhat more convenient terms, with *roles* objects occupy *relative to* utterances of them. For example, 'I' is associated with the utterance-relative role of *being the speaker*. A given utterance *u* of 'I' refers to the speaker of *u*.

Here are some plausible utterance-relative roles for familiar indexicals:

An utterance *u* of 'today' refers to the day on which *u* occurs
An utterance *u* of 'you' refers to the person to whom the speaker of *u* is speaking.
An utterance *u* of 'yesterday' refers to the day before the day on which *u* occurs

'This,' and 'That,' are demonstrative pronouns when used alone, and demonstrative adjectives in phrases like 'this pencil' and 'that pencil sharpener.' Given our working definition of indexicals, demonstratives are a species of indexicals. In the paradigm case, the referent of 'this' or 'that' will be an object that the speaker is attending to and to which he is directing the attention of his audience. This is not a matter of convention, but involves a real relation between the utterance and the object entails the perceptions and intentions of the speaker. The fact that in English 'this' is the word that has been assigned that role is, however, a matter of convention.

This,' 'that,' I,' and 'you' are all pronouns. But not all pronouns are indexicals, and not all indexicals are pronouns. 'Today,' 'tomorrow,' and 'yesterday,' for example, are adverbs. The pronouns 'he and 'she' can be used indexically, as demonstratives, or – unhelpfully – they can also be used in ways that do not clearly fall under the definition of indexicals. In,

Elwood is a man so he is mortal

the word 'he' is used with 'Elwood' as antecedent, and refers to Elwood because its antecedent does. In

Every young man thinks he is immortal

the 'he' does not refer to anything, but functions more or less as a bound variable.

Now consider this statement,

(5) Harry likes to drink at the local bar

The word 'local' is connected with the relation *x is near y*. Here the description 'the local bar' identifies the bar Harry likes to drink at as one that is near something. But what?

Contrast this with

(6) Harry likes to drink at the bar that is near his house

In (6), the parameter of being the thing the bar is near is represented by an explicit argument place, filled with the noun phrase, 'his house.' In (5) there is no explicit argument place for that parameter. I call expression like 'local,' which leaves an essential parameter of the relation that they express unarticulated, *role* expressions. Suppose by 'the local bar' in (5) we mean the bar near Harry's house, and that bar is McGinty's. Then McGinty's bar is the occupant of the role, and Harry's house is the anchor for the role, the thing relative to which the occupant has the associated relation.

Mandatory indexicals are role expressions, where the anchor is fixed by the rules of language to be the utterance of the expression itself. The referent of an utterance of 'I' is the speaker of that utterance, but there is no argument role in 'I' to indicate which utterance is in question. The fact that it is the very utterance whose reference is in question is fixed by the rules of language, so it need not be articulated. There is no option.

In the case of 'local,' there is an option. (5) can be read as saying that Harry likes to drink at the bar that is near the place of the utterance. If we were driving through a neighborhood in which there is just one bar, and Harry, although he lives miles away, is very fond of that bar, (5) would be appropriate. Or it could be the bar that is local to the neighborhood about which we are seeing a television documentary. Or it could be the bar that is near Harry's house. The rules of language do not fix which location is the anchor for the role.

Such *optional* indexicals are role-expressions that can have anchors mediated by previous parts of the discourse (or later parts, in some cases). So, we can say

(7) Harry wants to move to a town on the beach. He likes the local bar

and mean the bar that is local relative to the town on the beach to which Harry wants to move. We can also bind the anchor, as in

(8) Everywhere that Harry lives, he likes the local bar

These things are not possible with indexicals where the anchor is fixed as the utterance. For example,

(9) Harry made an important utterance. He said I am getting old

cannot be read as saying that Harry said that the speaker of the utterance that Harry made is getting old. It means that Harry said that the speaker of (9) is getting old.

All the senior citizens said I am getting old

means they all of the senior citizens said that a certain person, the speaker of (10), is getting old, not that each one said the she (herself) or he (himself) was getting old.

Demonstrative pronouns, such as 'that man,' are more flexible. We can say

(11) Elwood was upset by a cabinet member. He's very suspicious of that man.

The pronouns 'he' and 'she' can be used demonstratively, anaphorically or as bound variables:

(12) Harry thinks she [pointing] is intelligent
(13) Harry has a new neighbor. He thinks she is intelligent.
(14) Whenever Harry has a new boss, he thinks she is intelligent

These phenomena suggest that a proper understanding of role words should help provide a unified treatment of indexicality, anaphora, and quantification.[4] In this essay, however, we will focus on the mandatory indexicals and on the indexical uses of optional indexicals.

Indexicality, then, is not a syntactic category. Indexicality is a semantic category, having to do with meaning, reference, and truth. It is also a pragmatic phenomenon, in both senses of the term. In one sense of 'pragmatics,' it means aspects of meaning that depend on the properties of particular users; this is how Montague used the word in his essay "Pragmatics" (Montague 1968). Pragmatics is now usually conceived as the study of how speakers use their utterances to achieve goals of communication. Although what makes an expression an indexical is a semantic issue, the ways we use indexicals and the reasons they are important can only be understood within a pragmatic account. Pragmatics is also needed to help us understand how our use of indexicals adapts to changes in the basic communicative situation brought about by technologies of various sorts.

In *Reference and reflexivity*, I classified indexicals with a two-fold distinction:

- Does designation depend on narrow or wide context? Narrow indexicals depend only on the constitutive facts of an utterance: speaker, time, and place. Wide indexicals depend on other facts.
- Is designation 'automatic' given meaning and public contextual facts, or does it depend in part on the intentions of the speaker? (I called these indexicals "intentional" in *Reference and Reflexivity*, but here I call them "discretionary.")

Of course, with all expressions the intention to use them with their ordinary meanings is relevant. However, it seems that with a word like 'I,' no further intention is relevant to determining the referent. If I am speaking English, 'I' refers to me when I use it (although we will consider a possible counterexample below in section 5).

The indexicals 'now' and 'here' seem at first glance as automatic as 'I.' But with 'now' there is a question of how long an interval of time as counted as the present moment; with 'here' the is a question of how much of the surrounding territory is counted as the place of utterance. It seems there is a bit of additional intention that is at least possibly relevant to the determination of reference. Hence, these two are demoted to Discretionary/Narrow.

The ordinary demonstratives and third-person pronouns are clustered in the wide/discretionary cell. The reference of a use of 'that man,' for example, is not determined *merely* by the meaning of the expression, and the speaker, time and place of utterance. Wider facts are relevant. I see the determination of reference coming in two stages. First, which objects of the appropriate sort are *salient?* This is not a matter of the speaker's intention. Second, given the set of salient objects, the speaker's *directing intention* chooses among them.

Suppose I say "that man takes the money," to you, pointing at the man behind the cash register. I do this in order to get you to believe that that man is the one to pay, so that you will pay him, so that we can leave the restaurant. What links my speech to the rest of my plan is the belief that the man I am attending to and directing your attention toward is the cashier. My plan of reference is that I intend to refer to the fellow I see behind the counter, and *thereby* refer to the cashier, for I believe that man behind the counter is the cashier. It is the lower level intention, to refer to the man behind the counter, and not the higher level intention, to refer to the cashier, that is the direction intention, the one determinative in fixing the reference of my utterance. If my belief is wrong, I will succeed in referring to the man behind the counter, but I will *not* succeed in referring to the cashier. If the man behind the counter is cleaning the cash register, while the cashier stands outside having a smoke, I will have said something false about the man behind the cash register, not something true about the cashier. It is such directing intentions that are relevant in determining reference, not the various higher level intentions to refer that we hope to fulfill by carrying out the directing intention. The idea and the term "directing intention" are due to Kaplan (1989).

Let us now turn to the words in the Automatic/Wide cell. One might say, holding one's hands a foot apart, "The bass was yea big," or holding one's hand a couple of feet off the ground, "Her dog was yea big." 'Yea' is wide, because its

Table 17.1 Types of indexicals

	Narrow	Wide
Automatic	I	yea, dthat(α)
Discretionary	Now, here	That, this man, there, he, she, it

reference depends on how one holds one's hands as one says it. It is automatic, because the distance between one's outstretched hands is the distance to which one refers with 'yea,' whether one manages to match the distance one has in mind or not. One may have intended to hold one's hands a foot apart but have actually held them fourteen inches apart. Then the statement is false if the fish in question was exactly twelve inches long. It is not one's intention, but the distance one holds one's hands apart, that determines the semantic value of 'yea.'

On this account, then, the demonstration that accompanies a use of 'yea' differs in significance somewhat from one that accompanies of use of 'that man.' The demonstration is essential to the use of 'yea'; an aspect of the demonstration itself, the distance between the hands, is referred to. The speaker's intention is relevant in his decision to exhibit the length to which he intends to refer, but once he had done that, there is no discretion left. In the case of 'that man,' however, the demonstration is not essential. The demonstration helps to direct our attention to the object referred to, rather than determining which object is referred to.

Kaplan's invented demonstrative 'dthat(α)' *automatically* refers to the object which *in fact* fits the description α. It is *wide* because any sort of fact at all, and not just facts about who is speaking, where, and when, can be incorporated into α. It is automatic, because one refers to the person or thing that fits the description, whatever or whomever one might have wanted to refer to. For example, if you think Jefferson was the first President, and intended to refer to Jefferson by uttering "Dthat (the first President of the U.S.), you would have failed. You would have referred to Washington instead.

Later on, we will amend this table.

3 Reflexivity and Direct Reference

The reflexive-referential theory treats *meaning* as a property of expression types, and *content* as a property of utterances. Meanings are basically rules that determine the content for specific utterances. The contents of statements are propositions that capture their truth-conditions; the contents of subsentential expressions are the semantic values an utterance of them contributes to the contents of the statements of which they are parts. The English sentence 'I am happy' has the same meaning each time it is used (setting ambiguities, subtleties and odd uses aside). Different utterances of it have different contents, however, since the truth of those utterances depends on different people being happy at different times.

On the reflexive-referential theory, utterances have a variety of contents, the most important of which are *reflexive contents* and *referential contents*. The referential contents of utterances of sentences containing names, indexicals, and demonstratives will be just those assigned by standard referential theories.[5] The referential content of

(15) I am happy,

uttered by me, is the proposition

(16) that **John Perry** is happy.

I use bold face to identify the particular constituents of the proposition. The constituents are the subject matter, the things the proposition is about. (16) is what Kaplan calls a "singular proposition" true in worlds in which I am happy—a proposition with me as a constituent, rather than any identifying condition or "mode of presentation" of me.

I use italics to indicate that an identifying property, rather than the object that fits it, is the constituent of a proposition. So the proposition

(17) That ***the speaker of*** **(15)** is happy

has an identifying condition as constituent; it is what Kaplan calls a "general proposition" and is true in worlds in which whoever uttered (15) in the world is happy there. (18), on the other hand, is the same proposition as (16), a singular proposition with me as constituent:

(18) That **the speaker of (15)** is happy.

Kaplan and others use various arguments to show that (16) and not (17) is the proposition expressed by my utterance of (15); (16) is *what I said* (Kaplan 1977). I agree with this. I call (16) the *referential content* or *subject matter content* or *official content* of (15). (16) captures what are sometimes called the *counterfactual truth conditions*. These are the worlds in which the official content is true; worlds such, if they were actual, the proposition would be true. It is also crucial to recognize (17) as the *reflexive* content of (15). I call it "reflexive" simply because it is a condition on (15) *itself*. (15) is not what (15) is *about*; it is not part of the subject matter of (15). (17) is not the counterfactual truth condition of (15); the proposition expressed by (15) is true in worlds in which (15) itself does not occur. Still, (15) will be true if (17) is true, and vice versa. I claim that the reflexive content helps us understand the reasoning that motivates the production of utterances, and the reasoning that is involved in their interpretation.

4 Using Indexicals

Epistemic and pragmatic roles

When an object plays an *epistemic role* in our lives we have ways of finding out about it. If I am holding something in my hand, I can look at it, feel it, smell it, and so forth to get information about it. If I am standing in front of someone, I can open my eyes and look, ask questions, walk forward and touch, and so forth. Some epistemic roles are utterance-mediated. If you are telling me about your mother, I can ask you questions to find out more about her. I am in an utterance-mediated epistemic relation to your mother. Objects which are playing epistemic roles in our lives, or could easily do so if our attention were suitably directed, are *salient*.

If something plays a *pragmatic role* in our lives, then we can affect it, or use it to affect other things. We can do things with it and to it. Some pragmatic roles are utterance-mediated. I can thank your mother by asking you to convey my thanks to her next time you see her. Some of our examples of epistemic relations were also pragmatic relations. If I am holding something I can squeeze it, throw it, give it to you, and so forth. If I am standing in front of you, I can shove you, annoy you, or startle you. These are *epistemic-pragmatic* relations. However, not all epistemic relations are pragmatic relations. If am looking through binoculars pointed at a bird on a distant tree, I can find out about it, but cannot do much to have an effect on it. If I control big guns on a battleship I can have a devastating effect on the distant countryside, but having a big gun doesn't give me a way of knowing about the countryside I plan to shell.

Being at is also an epistemic-pragmatic relation. I can typically find out about the place I am at by looking around, and can effect it in countless ways. *Being at* may not be a causal relation, but it makes possible many sorts of causal interactions. I can plow the field I am in, paint the room I am in, and so forth.

We live in a world in which technology has created epistemic-pragmatic relations of all kinds. You and I are on opposite sides of the world, sitting at computers that are in turn hooked into the internet. We can exchange email; I can use the relation of *being on computers hooked to the internet* to find out about you and to affect you. I can also stand in an epistemic relation to objects around you, if you are telling me about them via email or your webcam. In addition, I may stand in a pragmatic relation to them, if you are willing to follow my instructions about what to do with them. I will return to this theme below. But for now, let us bracket technology. Let us think about the relatively simple and direct epistemic and pragmatic relations, as they were at the time our basic set of indexicals were developed—long before the internet. My hypothesis is that these indexicals are associated with roles that serve conversational purposes by directing the hearer to a second channel of information about their referent.

These epistemic and pragmatic roles are occupied by objects (including times, places, and other people) relative to knower/agents. Among these objects are tokens. When I speak, I typically create the token I use, and directly affect its salient properties. I can also hear it as I speak it. I have an epistemic-pragmatic relation with the token. My listeners have an epistemic relation to it; by hearing it or seeing it, they can determine several of its properties. By standing in this epistemic relation to the token they stand in epistemic relations to other things: the utterance that produced it, the mind behind that utterance, and the object the mind is referring to. Being in a conversation puts us in a variety of epistemic-pragmatic relations with one another by putting us in an epistemic-pragmatic relation to our own tokens and an epistemic relation to the tokens produced by other conversants. If I am using the tokens I produce to tell you about things, then I have opened an *utterance-mediated* epistemic route to those objects for you, a way for you to get information about them.

Suppose you are looking across the room, where there is a man. You are in an epistemic relation to the man; you can find out more about him by looking, walking over to him and asking him questions, and so forth. You are also in a pragmatic relation to him. For example, you could ask him to give you a cigar, or play the cello for you. You do not do these things, although you like cigars and listening to the cello, since you have no idea whether the man smokes cigars or plays the cello. This is not something you can find out by looking closely—at least, not unless you are Sherlock Holmes.

Now I point to the man and say to you, "That man plays the cello." You hear the token I produce. If you know English, you will know:

> The utterance that produced this token is true IFF the person the speaker attends to ands toand seeks to draw my attention to, plays the cello.

Note that you are at this point in a utterance-mediated epistemic relation to the man. This is a rather complicated relation:

- You are in an epistemic relation to the token I produce, which you can hear.
- You are thereby in an epistemic relation to my utterance, since you can infer its properties from the properties of the token you hear; you know the utterance is true iff the person the speaker refers to plays the cello.
- You are thereby in an epistemic relation to my beliefs about the referent, as expressed by the rest of the sentence or clause.
- You are thereby in an epistemic relation to the man I am talking about, since you can infer properties of his from my beliefs, assuming they are accurate.

This utterance-mediated channel is the *first* channel my utterance opens. It would have been opened, even if I had used 'he' unhelpfully rather than 'that man' demonstratively. My use of the demonstrative and my demonstration open up a *second* channel of information about the man I am talking about. It tells you that by finding the man to whom I am attending, directing your attention towards him, and looking at him, you can find out *more* stuff about the very same person you are learning about by listening to me. When you combine the information from both channels, you learn that there is someone you can easily walk over to (knowledge obtained by looking at him), who does in fact play the cello (knowledge obtained by listening to me). So you do this.

My sentence, then, opens up for you an utterance-mediated channel of information about an individual for whom you have, or easily can have, another channel of information that is *not* mediated by the utterance. *The indexical indicates what that other channel is*—or at least provides a first step that makes it easy to find.

Suppose you are looking for the teacher of philosophy 10. I say, "JP is the teacher of philosophy 10." This opens up a channel of information. The person you want to know more about is the person the speaker is referring to. You can ask me more about "JP." If instead I say "I am the teacher of philosophy 10," I open up this channel. You can ask me about the person I am talking about. But I also open up a different channel. The person in question is not only the person the speaker is talking about, but also the speaker himself. You can hear what I have to say about myself, but in addition you can look at me and draw some of your own conclusions.

Honoring roles

Which roles do we assign to indexicals? There are countless relations that objects stand to us, that are relevant to our thought and action. It is clear that only a small portion of these is honored by having an indexical assigned to express it. For example, the movements required to pick up a coffee mug that is fifteen inches from me are quite different than those required to pick up a coffee mug twenty inches from me. I need to stretch my arm in a different way, and perhaps lean forward for the one that is further away. I know what to do because the situations will look slightly differently to me. Here are two roles that coffee mugs can play, then, that are cognitively different, typically connected to somewhat different perceptions and somewhat different actions. However, we have no indexicals to capture the difference. Our ways of talking are more coarse-grained than our ways of thinking. I would call either of the cups "that cup" if I was just looking at it, and "this cup" once I reached for it and picked it up.

The roles that we have honored by having indexicals and demonstratives assigned to convey them are those that are useful in opening useful supplementary channels of information a variety of recurring conversational situations. Let us look at some examples.

Suppose that Elwood is sitting on the sofa. Mel, in the kitchen, looking out at the people in the living room, asks if anyone would like a beer. Elwood says, "I'd like a beer." This utterance is well designed to achieve his goal in uttering it, to get a beer from Mel. It puts only a modest cognitive burden on Mel. He just needs to see who is speaking, and he will know to whom he should give the beer. He does not need to know much about Elwood, and in particular he does not need to know Elwood's name. Here the utterance-mediated channel gives Mel the information that a certain man, the one Elwood is referring to, wants a beer. The 'I' constrains the referent to be the speaker. This is useful because Mel can see the speaker. He visually learns where the speaker is sitting and what he looks like. Combining the information from the two channels, Mel knows that the person who looks a certain way and is sitting in a certain place wants a beer.

Suppose Elwood had said, "Elwood would like a beer." If Mel had not known who Elwood was, he would not know to bring the speaker a beer. The name would be as useless as 'he' used unhelpfully. Mel might naturally ask, "Who is Elwood?" and, when Elwood says, "I am," feel somewhat put upon. The word 'I' is perfectly suited in this case to link Mel's ideas of *is a person who wants a drink* and his perceptual ideas of *people in the living room*. It would be irritating if Elwood used his name rather than this ready-made device. It would probably be taken to suggest self-importance, as if everyone was expected to know who Elwood was. In both cases,

(19) I'd like a beer
(20) Elwood would like a beer

Mel learns that the person the speaker refers to would like a beer. In both cases, Mel sees where the speaker is located. With the first utterance, Elwood uses the indexical 'I' to *coordinate* these two sources of information. Consequently, Mel knows where to bring the beer. In the second case, Elwood does not coordinate. He leaves Mel in the dark as to whom he is referring, and so leaves Mel in the dark as to whom he should bring the beer. Elwood is not giving enough information.

If we assume Elwood is being helpful, this generates the implicature that we already have the missing information; that Mel knows who Elwood refers to with 'Elwood.' That is why if Elwood says the second thing it sounds pretentious; it implicates that everyone can be expected to know his name. Even when Charles DeGaulle used to refer to himself as "DeGaulle," it sounded pretentious, although he was probably correct, that virtually everyone knew who he was.

Now suppose that Elwood is not seated where Mel can see him, but is down the hall and around the corner. He shouts, "I'd like a beer." In this case, the second channel of information may not be very helpful. It will depend on how much Mel can get out of the sound of Elwood's voice. Maybe he knows Elwood and will recognize his voice; then it is helpful. Maybe the acoustics of the house are such that Mel can pinpoint the beer-needer's location just from the perceived direction of the request. In many cases, it will not be helpful. The point is that it isn't simply the knowledge that *the speaker* is the referent that is crucial, it is the nature of the additional channels, such as *being the person I'm looking at*, that this knowledge will allow the hearer to open that is important. In cases where the hearer can take no step beyond *the speaker*, the use of 'I' will usually be inappropriate.

Later, Mel has delivered the beers and the party starts in earnest. Mel and Elwood are talking. Elwood introduces himself, "I'm Elwood," and extends his hand. As a result, Mel learns Elwood's name. Notice that this simple transaction would not work without the indexical. "Elwood is Elwood" would not do the job. It provides no way for Mel to connect the person occupying one epistemic role (the person in front of him he is talking to) with the occupant of the other (the person that the person he is talking to is talking about). When he hears "I'm Elwood," Mel knows that for the statement to be true the speaker must be named Elwood. He knows that the speaker is the person talking to him, whose lips he sees move in cadence with the sounds. He learns that the person speaking to him, whose looks he can focus on and associate with the name, has the name 'Elwood.'

Other indexicals: *this, that, he, she, here, today*, etc. all are useful in common conversational situations. In each case, their utility consists in linking the object talked about with the occupant of another epistemic and/or pragmatic relation relative to the hearer or reader.

The skilled user of language keeps in mind the goal of getting ideas appropriately linked to one another in a listener's or reader's head, so she will have the desired thoughts. One wants the listener to think of the object one is talking about in a certain way, that will connect with what one wants them to know and to do with this knowledge. To do this, one must put some thought into the cognitive burden one is assuming, and whether it is likely to be met. The skilled leaver of messages on answering machines, for example, will think about whether it will be obvious to the person playing the machine when the call was made; if not, he will not express important information in terms that presuppose this knowledge, like 'today' and 'tomorrow.'

5 Tokens and Technology

Primordial conversation was face to face. This meant the time of utterance and the time of perception of the uttered token were, for all practical purposes, the same. The location of the speaker and the location of the hearer, on the other hand, could be significantly different. There was a natural dichotomy between 'here' and 'there,' one for the speaker's position, the other for the hearer's. Not so with 'now.' The contrast between 'now' and 'then' would not be between speaker's time and hearer's time, but time of utterance and some other salient time. Speakers and hearers share their nows, but not always their heres. Shouting and smoke signals allow communication at considerable distance, and the telephone takes that further—but the time of utterance and time of perception remain the same. With written language, however, the times of utterance and token-perception can be distant from one another. Copying allows multiple listeners at different times and places to perceive the same token; printing magnifies this effect. Email and the rest of the internet push the envelope in all of these directions.

Must I be here now?

The way we use indexicals has adapted to the situations technology makes possible. As we observed above, Kaplan argued in the late 1970s that 'I am not here now' could not be used truly, in any context. Given the rules,

(21) An utterance *u* of 'I' refers to the speaker of *u*

(22) An utterance *u* of 'here' refers to the position of *u*
(23) An utterance *u* of 'now' refers to the time at which *u* occurs.

it seems that Kaplan has a strong case. Leave aside the issue whether the certainty is a matter of logic, or a matter of a necessary truth about utterances, or simply a very well entrenched truth about utterances. Whatever the exact reason, it seems that if these rules are correct, no utterance of 'I am not here now' will be true.[6]

By the 1980s, answering machines had proliferated, and "I am not here now" became an oft-heard and easily understood and believed message. Is there something wrong with (21), (22), or (23)? Or something wrong with the 1980s?

One might suppose that it has become conventional to refer to the time of listening to the recording with 'now.' However, messages like

(24) I've got to leave now. I won't be here when you hear this recording. I'll try to call you back later.

in which the use of 'now' refers to the time of utterance, are permissible and intelligible. It seems to be the speaker's choice, whether to use 'now' for the time the recording is made or the time it is heard.

Stephano Predelli (1998a) argues that the phenomenon in question predates answering machines. Only the technology of leaving notes is required. Predelli's character Jones has to flee unexpectedly; he leaves his wife a note:

[P] As you can see, I am not at home now. If you hurry, you'll catch the evening flight Los Cabos. Meet me in six hours at the Hotel Cabo Real.

The token produced in this case is the note. Jones has a plan, that his wife see the note when she returns at 5 p.m. He uses 'now' to refer to that time, not to the time when he writes the note.

Again, we cannot simply suppose that there is a convention with notes to use 'now' for the time of token-perception. Jones could have written:

[P'] I'm leaving now for Los Cabos. I'll have been gone for a long time by the time you read this when you get home. If you hurry, you'll catch the evening flight to Los Cabos, and can meet me by 11.

If Jones had written P,' the use of 'now' would have referred to the time he left the note. It would be his intention to have his wife understand it as so referring that would be crucial.

Predelli advocates adding parameters for the intended place and time, which are the values for 'here' and 'now,' to Kaplan's contexts. The values may be the same as the time and place of utterance, but need not be. Jones intends to refer with 'now' to 5 p.m., the time he expects his wife to see the note. According to Predelli, this intention rules; had Jones' wife arrived home early or late, that wouldn't have made the 'now' refer to her time of arrival; it would have still referred to 5 p.m., when Jones expected her to read it. I am not so sure about this.[7]

Here is a somewhat different approach. In cases in which there is a separation between the time of utterance and time of token-perception, both times may be relevant and/or salient, and directing intentions may have a role. The note-writer or voice-recorder may have a directing intention that determines to which of two or more salient times he refers to with 'now,' as and ordinary speaker might with 'that man.'

The function of 'now' is to help the listener establish a separate information channel for the time the speaker refers to. In the face to face case the listener simply looks and listens at what is going on around him. If we are in a face to face conversation, and you say, "It's time for the meeting now," I do not need to examine your demonstrations or eye-gaze to see to which time you refer. I do not even need to make sure it is me to whom you are talking, as I might with 'you' or 'there.' The time to which you refer is the time that is present for both of us, and all within earshot, the time of utterance and the time of perception.

Writing changes this. You write me a postcard,

(25) Made it to Tokyo. I'm now very tired. Call you when I get back.

The simple techniques appropriate in face to face communication won't work. I need to know when you wrote the card. Perhaps I infer from the postmark that you probably wrote it last Tuesday; then I can find out about last Tuesday in all the usual ways—checking Wednesday's newspaper, for example. If I do not know when you wrote the postcard, 'now' is a pretty unhelpful indexical.

When the time of utterance is separated from the time of token-perception, the usual epistemic techniques associated with 'now'—listening and looking to what is going on—are techniques for the audience to find out more about the time of perception, but not the time of utterance.[8] If the time of token-perception is relevant, it will be a live candidate for the referent of 'now.'

From this point of view, the Jones case goes as follows. Jones expects his wife to read the note, and assumes that the time at which she reads it will be salient to her: it will be a live candidate to be the time he is referring to with his use of 'now.' It is relevant to what he wants to tell her, for it is the time relative to which the plane is leaving soon. He has a directing intention to refer to that time. He thinks she will see it at 5 p.m., so he thinks *by* referring the time when she reads the note, he will refer to 5 p.m. If his belief is wrong, and she reads it at 3:30 p.m., then he refers to 3:30 p.m.

Suppose the note contains the sentence

[P''] You must leave right now to catch the plane

This what actually happens. His wife sees the note at 3:30. She hurries to catch the evening flight, but finds she has to wait at the airport, having arrived an hour and a half before Jones thought she would. After they meet, she criticizes him, for saying something false, that she had to leave at 3:30 to catch the plane. According to the present approach, she is right, he did say that. On Predelli's analysis, he did not say that; he said that she had to leave at 5 p.m. to catch the plane, and she misunderstood.

Suppose Jones' expectation that his wife would return at 5 and see the note then was based on her usual schedule. Then she should have known when he expected her to read it, and he could reply that she should have figured out what he *meant* to say. It does not seem to me, however, that he could claim to have actually *said* it. The term 'said' is to a certain extent a *forensic* term; it has to do with what message people can be held responsible for conveying. At the same time, the hearer is responsible for using a bit of common sense, in figuring out what the speaker is trying to say especially when the speaker is obviously confused. Jones' wife has a legitimate complaint against Jones, but he has a legitimate comeback. A normal married couple.

On Varol Akman's view the word 'I' has the speaker as a *default* reference, but in certain circumstances 'I' can refer to others. I'll express some doubts about his case for the word 'I' below, but here I want to adapt his suggestion to the word 'now.' The default value for a use of 'now' is the time of utterance, in the sense that the speaker can always use 'now' to refer to the time of utterance. He can use 'now' to refer to the time of token-perception when it is salient and relevant.

Where is here, anyway?

Let us think a bit about "here." Suppose Jones's situation was more dangerous than Predelli suggests. He called his wife and told her not to go home at all but to look through a telescope from a neighbor's house at a note he would leave on the refrigerator with instructions about what she should do. The note begins

I am not here now . . .

In the situation, the 'here' seems to refer to the house where the note is, the 'now' to the expected time of perception of it. It does not seem possible to read 'here' as referring to the neighbor's house, the place of perception. If the note continued:

I am hiding there, in the front-hall closet . . .

Mrs. Jones would look for her husband in the neighbor's front-hall closet. She would take 'here' to refer to the place where the note was, 'there' to the place where she was.[9]

These considerations seem to me to weigh in favor of my account. Both the place of the note and the place from which Mrs. Jones perceives the note are salient. However, it seems Jones should refer to the first as "here" and the second as ‘‘there.’’ The reason is that the contrast between ‘here’ as the place of the speaker and ‘there’ as the place of the listener is well established, as is the difference in epistemic techniques. If I say, ‘‘look over there’’ without any further indication of a place, you will take me to have told you to look around the part of the world you are at. If I say, ‘‘it’s over here’’ you would look for the desired object close to me.

When we return to the telephone answering machine, things get rather murky, however. Sometimes answering machines are located near the telephones, calls to which they pick up after a few rings. Sometimes, however, the answering machines are at a central location, perhaps a telephone company office, or the communications department of a corporation. When the message says, ‘‘I am not here now’’ the ‘‘here’’ seems to refer to the place where the telephone is, whether or not the answering machine is there.

The use of ‘here,’ ‘there,’ ‘now,’ and ‘then’ as contrastive demonstratives is also confusing. This is most clear when there is some sort of representation we are using to discuss events. With appropriate demonstrations to a map of an intersection in a courtroom, I can say, ‘‘I was stopped here. He backed out of his garage there and . . . ’’ I can use ‘now’ and ‘then’ similarly with a chronology of events. What, if anything, do these uses have to do with the time and place of utterance? ‘Now’ and ‘then,’ seem to amount to ‘this time’ and ‘that time,’‘here’ and ‘there’ to ‘this place’ and ‘that place.’ I’ll, reluctantly treat these uses as separate senses, subscripted with ‘D’ for ‘demonstrative,’ until a clearer vision allows a unifying account.

Here then is a revised table.

There seems to be a steady drift away from the upper left corner. Only ‘I’ is left as automatic and narrow. Can it hold that ground?

Must I be me?

Varol Akman suggests that ‘I’ is more flexible than its lonely position in the narrow and automatic cell indicates.[10] Akman imagines an ill Yeltsin, looking at the man who has been serving as his double asking ‘‘How am I doing today?’’ If we take the ‘‘I’’ to refer to the double, then we seem to have a case where ‘I’ does not refer to

Table 17.2 Types of indexicals

	Narrow (speaker and time and place of utterance are only relevant facts)	Less narrow (time and place of token-perception are also relevant)	Wide
Automatic	I		Yea, dthat(α)
Discretionary		Now, then, here there	That, this man, he, she, it, now_D, $then_D$, $here_D$, $there_D$

the speaker. Taken this way, Yeltsin would be asking the double how the double was doing; was he hot? tired? bored with pretending to be Yeltsin? and so forth.

Suppose Yeltsin was at a conference with some other big shots, all of who had doubles. Stopping by the double buffet for a drink, Yeltsin asks a series of questions, addressing each to the appropriate double: How is Bill Clinton today? How is François Mitterrand today? How is John Major today? How am I today? It seems to me that to understand what is going on we don't need defeasible interpretations for the names, but rather what W. V. Quine called "deferred ostension" (1969). There is a "proxy-function" ρ. By referring to X one identifies, and talks about, the ρ of X. One car park attendant says to another, "Here comes the Porsche," meaning the owner of the Porsche is coming to get his car. Thus, Yeltsin asks about Bill Clinton's double, the person in front of him, when he says, "How is Bill Clinton today." Yeltsin can pretend to be taken in, by referring to the doubles the way the deceived commoners are expected to refer to them. The joke comes when he uses 'I,' when the deferred ostension still works, but the pretense runs aground; he cannot refer to his double as 'I' without betraying that he realizes that it is a double, and not the real Yeltsin. The case is interesting, and the default/defeasible is plausible in the case of deferred ostension. However, the default is that the person referred to is the person talked about or asked about, the propositional constituent of the what is said or what is asked. If we approach Akman's example this way, we can leave 'I' its honored place in the chart.

Notes

1 I use double quotes for quoting language and thought, and as scare-quotes; I use single quotes for mentioning expressions.

2 As I mentioned above, I use "real relation" for one not mediated by utterances. As a matter of convenience, not scholarship, I'll take Peirce to mean something by "real relation" that implies my sense.

3 Reichenbach's own approach was a bit different, as was what he meant by "reflexivity." I am stating the truth-conditions of (1) and (2), not trying to produce synonymous translations for them. Reichenbach wanted to produce a synonymous symbolic formula. For this purpose he introduced one all-purpose indexical, τ, which means roughly "this very token." For (1) this would have been something like

(1R) $\exists x$ x is the speaker of τ & x has been to Paris

(1R), like (1) and (2), is token sensitive. The difference is that Reichenbach has isolated the token-sensitivity to a single word, his invented token-reflexive τ.)

4 See ch. 18.

5 I use the term "referential" for theories that take names and/or indexicals to be "rigid designators" (Kripke), or to be "directly referential" (Kaplan) or to take statements containing such expressions to "say something about" the object referred to (Wettstein, Donnellan). The idea is that sentences using these terms will be about

the objects referred to, rather than any mode of presentation or identifying description of them.

6 For discussion of the status of "I am here now," see Predelli 1998b and Perry forthcoming.

7 I agreed with Predelli on this point in Perry, forthcoming.

8 Idea: salience as availability of independent channel of information.

9 Note also that in the case of the answering machine's message "I am not here now," it is the place where the telephone is or is expected to be, not the place where the answering machine is. Some answering machines are provided by telephone service providers, and are located miles from the phones they answer. I may use such an answering machine, and in addition, for some reason or other, disconnect my telephone and bring it with me when I take a short trip. So there is no telephone at my home, and the answering machine you hear is not there either. Still, it seems that "here" refers to my home, the location where the telephone you were trying to call was expected to be. If you called my office, where I had also removed the telephone, and received a message from the same central location for answering machines, the 'here' would refer to my office. The application of all of this to cell telephones is left to the reader.

10 V. Akman, "Context and the indexical 'I'," paper read at NASSLLI'02 Workshop on Cognition: Formal Models and Experimental Results, CSLI, Stanford, CA (30 June 2002).