

Reflexivity, Indexicality and Names

John Perry

[In *Direct Reference, Indexicality and Proposition Attitudes* edited by Wolfgang Kunne, Martin Anduschus, and Albert Newen. Stanford, CA: CSLI-Cambridge University Press, 1997.]

1 Introduction

Consider:

- (1) “I am a computer scientist.” (said by David Israel)
- (2) “David is a computer scientist.” (said by someone who is referring to David Israel with “David”)

It has been persuasively argued by David Kaplan and others that the proposition expressed by statements like (1) is a *singular proposition*, true in just those worlds in which a certain person, David Israel, is a computer scientist. Call this proposition *P*. The truth of this proposition does not require that the utterance (1) occur, or even that Israel has ever said anything at all. Marcus, Donnellan, Kripke and others have persuasively argued for a view of proper names that, put in Kaplan’s terms and applied to this example, implies that the proposition expressed by (2) is also simply *P*.¹

The thesis that expressions of a certain category (names, indexicals, demonstratives, pronouns, descriptions, etc.) are *referential*² holds that these expressions contribute the object to which they refer, rather than a mode of presentation of that object, to the propositions expressed by statements containing them.

The thesis that indexicals and names are referential creates the challenge of explaining the difference in *cognitive significance* between statements like (1) and (2), that express the same proposition [Wettstein, 1986]. The problem has two parts, which

¹[Marcus, 1961], [Donnellan, 1970],[Kripke, 1980]

²David Kaplan’s term is *directly referential*. Kaplan has a precise concept of “directness” in mind, but unless one is focusing on his exact words, the term “directly” is likely to suggest that there is no semantic mechanism intervening between the expression and its referent. This is pretty clearly not the case with indexicals, as Kaplan’s own analysis shows; it may be more plausible for proper names. Using terminology introduced in Section 3 we can say: Kaplan’s language suggests that directly referential terms *name*, but what he really says is simply that they *refer*.

I'll call "cognitive motivation" and "cognitive impact". The problem of cognitive motivation is that in many circumstances, the beliefs and desires that might lead a competent user of English to assert (1) differ from those that might lead one to assert (2). The problem of cognitive impact is that in many circumstances the beliefs and desires a competent listener would acquire (and would be expected to acquire) from hearing (1) differ from those she would acquire from hearing (2). If (1) and (2) express exactly the same proposition, why should this be so?

Let's look at our example in some detail. First note that differences in David's beliefs might lead him to assert (1) rather than (2) or vice versa. Suppose David has amnesia. He might have figured out what his profession was, but not have remembered his name. Then he would be in a position to assert (1) but not (2). Or he might have figured out that a certain person, David Israel, was a computer scientist, without realizing that he is that person. Then he would be in a position to assert (2) but not (1).

In the normal case, in which David does not have amnesia, differences in desire could motivate him to assert (1) rather than (2), or vice versa. In normal circumstances, if Israel wants to tell someone his profession, he will use (1) and not (2). Suppose, for example, that David has made some interesting points following a paper by a guest lecturer at CSLI. After the session he chats with the speaker, who seems unsure that he can trust what David says about programming languages. It would be natural for David to say (1), and odd for him to say (2). He wants to reassure his interlocutor about the reliability of what he says about programming languages. (1) will do that; (2) might not. I'll call this example Case 1.

There are easily imaginable circumstances in which David's desires would lead him to use (2) but not (1). Suppose a government bureaucrat has the job of identifying computer scientists, who will then be sent a letter offering a well-paid job building the Information Highway. The bureaucrat himself has a clear antipathy towards computer scientists. He reads David names from a list of people who work at SRI and asks if they were then or ever had been computer scientists. When he gets to "David Israel", Israel might just say (2), figuring that the information that *he* was the person in question was simply irrelevant and would lead the bureaucrat to be rude to him. Here David would not want his interlocutor to know that he, the person being talked to, was a computer scientist but he would want him to know the individual who is designated by the token of "David Israel" on the list is a computer scientist. He thinks that he will pass on that information with (2). I'll call this Case 2.

Given that David knows that he is David Israel, his choice between (1) and (2) is based on the effect he wants his utterance to have on the beliefs of his interlocutor—what I am calling its "cognitive impact". He thinks that if he utters (1), his interlocutor will realize that the person he is talking to is a computer scientist. If he utters (2), his interlocutor might not realize that, even if he believes what David says.

In fact Israel is well-known as a computer scientist, but looks and acts more like a philosopher. A competent speaker who hears (1) may well be skeptical about what he hears, even though he readily assents to (2), or even asserts it himself, having read articles by Israel. How is this difference possible, if the propositions expressed by (1) and (2) are the same?

My thesis in this paper is that to understand the difference in cognitive significance between (1) and (2), we need to acknowledge several different kinds of *content*. The paper develops a line of argument put forward in the later papers of [Perry, 1993]. In those papers a distinction is made between *reflexive* and *incremental* content, and it is argued that the former is relevant to cognitive significance. This distinction is most intuitive in the case of indexicals. In this paper I show how the basic idea can be extended to deal with proper names, without treating them as if they were indexicals. This requires examining the difference between reflexivity and indexicality, and explaining how the concept of reflexivity can be applied to expressions that are not indexicals. But I will start with indexicals, and then move to consider names.

2 Indexicals

(A) captures most of what a competent English speaker knows about the word “I”:

(A) A use u of “I” designates x iff x is the speaker of u .

From this, we see that (1), will be true iff

Q : There is an x such that i) x is the speaker of (1);
ii) x is a computer scientist.

Q is not the proposition expressed by (1); as we saw, that is P . The truth values of P and Q will agree in the actual world, but may disagree in other possible worlds. In the worlds in which Israel is a computer scientist, but does not utter (1), P is true and Q is false. This does not mean that Q is irrelevant to the cognitive significance of (1), however. Q gives us what a competent user of English who hears (1) will know, just on the basis of those facts, about its truth conditions.

Someone who knows English and hears (1), but isn’t in a position to see who said it, would still realize that Q captured the conditions under which it was true. Such a person might express this knowledge by saying “This utterance is true, iff *its* speaker is a computer scientist.” The truth conditions he associates with (1) are *reflexive* in that Q is about the utterance (1) *itself*. Because he doesn’t know who the speaker is, he cannot “detach” the truth conditions of (1) from (1) itself.

In contrast, someone who learns that Israel is the speaker of (1), can get at the truth conditions of (1) in a non-reflexive way: “This utterance is true iff Israel is a computer scientist.” I call this the *incremental* truth conditions of (1), *given* that Israel is the speaker. “Incremental” is used here in the sense of “additional”. Given that Israel is the speaker of u , what *more* has to be true for u to be true? Answer: Israel has to be a computer scientist. This is just the proposition P , the proposition expressed by both (1) and (2). The incremental truth conditions correspond to the ordinary philosophical concept of truth conditions; they are *what is said* by the speaker of the utterance. I’ll also call this “official content”.

The difference between Cases 1 and 2 above has to do with the *mode of presentation* that Israel wanted his hearer to associate with the belief that he was a computer scientist. In Case 1, Israel wanted the hearer to recognize that the person spoken about was also the person speaking. The plan of understanding that he has in mind goes roughly like this:

The hearer will perceive my utterance, and, as a competent speaker of English will recognize that it is true iff the person who makes it is a computer scientist (Q). He will also see that I, the person in front of him and talking to him, is the speaker of the utterance. So he will learn that the person in front of him is a computer scientist, and will realize he can learn more about this fascinating field by asking questions of this person.

In Case 2, when Israel does not want the hearer to recognize that he is speaking to a computer scientist, he does not refer to himself with the expression “I”, and so does not disclose to the hearer that the person with whom he is talking is a computer scientist.

We cannot explain all of this if we only have the official content of (1) to work with. Its official content is the singular proposition P , the same official content that (2) has. But that is not all we have to work with. We also have the reflexive content, the proposition Q . This allows us to construct an explanation of why David chose to utter (1) in Case 1 but not in Case 2.

One might object at this point as follows. The source of the original problem was the fact that according to referentialism we don’t have modes of presentation of David Israel but Israel himself in the propositions expressed by (1) and (2); the difference between them is lost: they both express the same singular proposition, P . But we cannot solve this by appealing to Q , for it is also a singular proposition, containing the utterance (1) as a constituent. Just as one can have different modes of presentation of David Israel, one can have different modes of presentation of (1). In Case 1, Israel thinks of (1) in the way one thinks of an utterance one is making, while the guest lecturer thinks of (1) as an utterance made by the person one is talking to.

Suppose the room in which David utters (1) is large and cavernous and gives rise to echoes that are so clear people often fail to recognize them as echoes. A few seconds after Israel utters (1) and the guest lecturer hears him, the guest lecturer hear the echo, “I am a computer scientist”. This gives the guest lecturer second mode of presentation of David’s utterance. The first time he heard it, he thought of it as (roughly) “this utterance being made by the person I am talking to.” The second time, when he heard the echo, he thought of it as “this strange utterance I am hearing I don’t know from where”.

These two hearings have quite a different cognitive impact on the guest lecturer. But the same utterance is involved both times, and hence the same reflexive content, the proposition Q with David’s utterance as a constituent. The guest lecturer has two modes of presentation of (1), corresponding to the two times he heard it. Since Q is a singular proposition with (1) itself as a constituent, rather than a mode of presentation, Q can’t account for the different impact the two hearings had. So Q , because it is a singular proposition, can’t solve our problem.

This objection fails. The fact that P is a singular proposition gives rise to the problem of saying how the cognitive motivation and impact of (1) and (2) differ, given that they both express P . To note that Q is the reflexive content of (1) but not of (2), allows us to understand how there can be a plan into which an utterance of (1) fits but an utterance of (2) does not. That's all we need to do. We could now construct an example in which two different utterances had Q as their content, say

α "There is an x such that i) x is the lecturer of this utterance and ii) x is a computer scientist," where "this utterance" refers to (1).

β "There is an x such that i) x is the speaker of (1) and ii) x is a computer scientist."

This would raise a problem analogous to the one raised by (1) and (2), and could be solved along the same lines. The reflexive content of α would be a singular proposition having the utterance α itself as a constituent, and not having (1) as a constituent, and this reflexive content could be used in explaining the motivation for uttering α and the impact of uttering α in cases where the motivation for uttering it or the impact of hearing it differs markedly from that of uttering or hearing β .

Still, we have only solved part of our problem. In case 2, David does want his hearer to realize that the person designated by "David Israel" on his list of SRI employees is a computer scientist. He manages to bring this impact about by using his own name when he responds. How does this work?

3 Names

Before I develop an account of names, let me make two important distinctions about designation. The first distinction has to do with what a term contributes to the official content of statements of which it is a part. Terms *refer* if they contribute the object they designate, rather than some mode of presentation of it, to official content, so that their official contents are singular propositions about the object designated. Terms *describe* if they contribute a mode of presentation, a condition on objects, that they incorporate. In the table below, referentialism about indexicals and names is reflected by entries in the “Refers” column. Descriptions are put in the “Describes” column, since we are not assuming referentialism with respect to them ³.

	Refers	Describes
Names	Proper names	???
Denotes	Indexicals	Descriptions

Table 1: Varieties of Designation

I use the words “denotes” and “names” for a different distinction. This has to do with the kind of conventions that are associated with terms of a given type. Terms *denote* if the conventions of language associate them with modes of presentation, with conditions on the object designated. This is clearly true of descriptions. For example, the rules of language do not associate any particular person with the description, “the tallest philosophy professor in America,” but merely a certain condition. The person who satisfies the condition is designated by the description, but one can know the meaning of description without having any idea who that is.

Indexicals also denote, in this sense. The conventions of English do not associate particular individuals with the words “I” and “you”, but conditions that individuals must satisfy to be the designata of uses of those words. So indexicals denote, as descriptions do. But indexicals do not describe. They refer. The individuals that meet the conditions, rather than the conditions themselves, are contributed to official content.

Terms *name*, on the other hand, if the conventions that secure designations for them directly associate them with objects, rather than conditions on objects. This distinction and terminology allows me to state my account of names in a way that sounds persuasively trivial: names name.

When a person or thing is assigned a name, a *permissive convention* is established: that name *may* be used to designate that person. When David Israel’s parents named him “David”, they established a convention that made it possible for people to designate their son with the name “David”. It did not preclude people from using “David” to designate other people.

³For this distinction see [Marti, forthcoming]; see also [Recanati, 1993] for wisdom about the issues.

When a name is used in a conversation or text to refer to a given person, the speaker is exploiting a permissive convention of this sort. A single name like “David” may be associated with hundreds of thousands of people by different permissive conventions. In the abstract, the problem of knowing which conventions are being exploited when one apprehends a token containing the word “David” are considerable. And when one sees something like “David was here” scrawled on a wall one may be completely clueless as to which of the millions of Davids is being designated. But usually various factors work to make the use of proper names a practical way of talking about things. I only know a small minority of the Davids that can be designated with “David”; the ones I know overlap in various fairly predictable ways with the ones known by people I regularly meet in various contexts; principles of charity dictate that I take my interlocutors to be designating Davids that might have, or might be taken to have, the properties that are being predicated of the David in question; and I can always just ask.

The role of context in resolving the issues of which naming conventions are being exploited is quite different from its role with indexicals. In the case of indexicals, the meaning of a given expression determines that certain specific contextual relationships to the utterance and utterer—who is speaking, or to whom, or when—determine designation. Different facts are relevant for different indexicals, and the meaning of the indexical determines which. Names don’t work like this. The difference between “David” and “Harold” is not that they are tied, by their meanings, to different relationships to the utterance or utterer. The role of context is simply to help us narrow down the possibilities for the permissive conventions that are being exploited.

If we have to give this phenomenon a familiar name, it would be “ambiguity”. The same name has many different meanings; as with ambiguous expressions, the role of context is to help us determine which meaning is relevant in a given use, rather than to supply a specific type of fact called for by the relevant meaning. There are many differences between the phenomenon in question and what we usually call “ambiguity”, however. Paradigm ambiguous expressions have only a few meanings, most of which are known to people who use the expression, or can be easily found by looking in a good dictionary. One can realistically aspire to knowing most of the meanings of many words. Many names have thousands of meanings—that is, there are thousands of individuals that they are used to designate, exploiting various permissive conventions. People who use a given name will be ignorant of the vast majority of its meanings, and it would be silly to aspire to know most of them. For help in discovering or narrowing the possibilities in particular cases one might use a phonebook, or even an encyclopedia, but not a dictionary.

So with names we have ambiguity, not indexicality. Nevertheless, we can define a useful concept of reflexive truth conditions for statements involving names.

Consider a case in which I see “David uses LISP” spray painted on a wall that I pass on my way to work in the morning. I have no idea which David is being designated. So I have no idea what proposition is expressed by the graffiti, no idea of what its official content might be. But I can say under what conditions the graffiti—call it *g*—is true⁴:

⁴I am sliding over the distinction between tokens and utterances here; what I perceive is a token that I take

- R*: There is a person x and a convention C such that
- i) C is exploited by g ;
 - ii) C permits one to designate x with “David”;
 - iii) x uses LISP.

R is not the proposition expressed by g —not g ’s official content. But it does give truth conditions for g . And R is reflexive; that is, it is a proposition about the very utterance of which it is the truth conditions.

Now let us suppose that g is actually a remark about David Israel, spray painted on the wall by a gang of teen-age admirers. In that case the content of g is a singular proposition about David Israel, and quite a different proposition than R . But R is a proposition we want to have available, as part of our explanation of the impact of the perception of this message has on me.

Suppose I am irritated with the graffiti, and want to find out who did it. I start with R and common sense. The people who wrote g were exploiting a convention that connected a particular David with their use of “David”. This David uses LISP (or at least might be accused of such a thing). And the graffiti writers in my neighborhood are in a position to exploit a convention that permits them to refer to him with “David”. This might allow me to figure out who the graffiti was about—I ask around the neighborhood to see if anyone knows a computer scientist named David who either uses LISP or would be shocked to be accused of such a thing. Once I track down David, I ask him who might be spray painting things on walls about his programming habits.

Now let’s return to Case 2. This is the case in which David utters (2) because he doesn’t want his interlocutor to know that he is talking to David Israel, although he doesn’t mind if he learns that the David Israel he is asking about is a computer scientist. We imagined that the interlocutor had a list of SRI employees, and he was going down the list.

First let’s ask ourselves what the bureaucrat learns when he sees the name “David Israel” on the list of people who work at SRI. The official content of the statement made by the presence of Israel’s name on the list is the singular proposition that he works at SRI. It would be misleading to say that this is what the interlocutor learns, however, since we are supposing he is merely a bureaucrat who has never heard of David Israel and has no concept of him except that he gains in virtue of seeing his name on the list. He has no way of thinking of Israel, that is, except as “the person designated by this use of ‘David Israel’,” which in terms of our account comes to “the person whom the conventions exploited by this use of ‘David Israel’ allow us to call ‘David Israel’”. We might then say that the bureaucrat does not fully grasp the meaning of the statement he is inspecting (which I’ll call h). He knows the conventions associated with “works at SRI” and he knows the type of convention associated with the use of “David Israel,” so he knows quite a bit. But he doesn’t know everything. What he knows is that h is true iff:

- S*: There is a person x and a convention C such that

to be a trace of and evidence for an utterance. See [Perry, forthcoming].

- i) C is exploited by h ;
- ii) C permits one to designate x with “David”;
- iii) x works at SRI.

Now the bureaucrat formulates his question: “What does David Israel do?” In asking this question, the bureaucrat exploits the convention that governs the token of “David Israel” on the list of SRI employees he is using. It is interesting that he can do this. He doesn’t know the convention *as* a convention for calling a certain person by a name, but only *as* the convention that is governing a certain use of the name. As far as he knows, he has never been presented with the person designated by that token except *as* the person presented by it. It may be surprising, if one has certain philosophical predispositions, that the bureaucrat can ask a question the official content of which has David Israel as a constituent, a question which is *about* David Israel, when he knows so little about him. But it seems scarcely deniable. The bureaucrat seeks to be a conduit, taking information about Israel from person in front of him, who he assumed to have some richer level of acquaintance and interaction with Israel than he himself has, and passing it on down the line until it will be used by someone who also has a richer level of acquaintance or can gain it, say by summoning Israel to an office for an interview.

To achieve this purpose, the bureaucrat needn’t know who Israel is, by any normal standards at least. What he does need to do is to make sure that the David Israel he has been asked to gather information about is the same one that the person in front of him is going to tell him about. He needs to make sure that the conventions governing his use of “David Israel” are the ones governing its use in the reply he gets. He might do that by asking, “Do you know a person named ‘David Israel’ who works here at SRI? Do you know more than one?” If David answers appropriately he goes on to ask, “And what does David Israel do?”

David replies with (2). The reflexive truth conditions of (2) are that

- T : There is a person x and a convention C such that
- i) C is exploited by (2);
 - ii) C permits one to designate x with “David”;
 - iii) x is a computer scientist.

David’s plan is that the bureaucrat, being a competent English speaker, will grasp T . He will assume that David is being helpful, and exploiting the same naming convention that he is—the same one the question on the form he is filling out exploits. This enables him to answer the question on his form. So David is helpful. But nothing in the transaction provides the interlocutor with information that David Israel is the person that is speaking to him. So, the level of reflexive content enables us to complete our story, and see how by using (2) David can expect to satisfy the combination of beliefs and desires that he had in Case 2.

4 Relative Truth Conditions

We have appealed to reflexive content to solve both parts of our problem, Q for (1) and T for (2):

Q: There is an x such that:
i) x is the speaker of (1);
ii) x is a computer scientist.

T: There is a person x and a convention C such that:
i) C is exploited by (2);
ii) C permits one to designate x with “David”;
iii) x is a computer scientist.

Although *Q* and *T* are both reflexive, in the sense that they are about the very utterance whose truth conditions they are, there is an important difference between them.

Earlier we contrasted *Q* with *P*, the proposition that David is a computer scientist. I called *P* the *incremental* truth conditions of (1). The idea is that we fix the contextual facts: that David is the speaker of (1). Now *given* this fact, what *else* has to be true for (1) to be true? What are additional or incremental truth conditions of (1)? Simply that David is a computer scientist. However, the binary distinction, “incremental vs. reflexive” is too simple. We need more flexible concepts.

Let’s step back a moment and think about the question, “What are the truth conditions of an utterance u ?” I claim that the form of this question should really be “Given that utterance u has characteristics C_1, \dots, C_n , what *else* has to be the case for it to be true?”

To explain what I am getting at here, let’s consider some variations on Case 1:

- a) The guest lecturer hears the words “I am a computer scientist,” but he is not sure which of the people in front of him said them.
- b) David mumbles and so the guest lecturer hears, “I am a [mumble-mumble].”
- c) David’s Boston accent is so bad that although the guest lecturer realized something has been asserted, he is not sure what the words were, or even what language they were in.

In these circumstances, does the speaker know the truth conditions of (1)? Not in the ordinary sense in which “the truth conditions” are used in philosophy. But in each case, there is a sense in which the guest lecturer knows the truth conditions, for in each case the speaker knows some facts about the utterance, and knows what additional conditions must be met for the utterance to be true:

- a) In this case, we fix the fact that (1) was uttered by a member of the group in front of the guest lecturer and the facts about which words were spoken in what language and what they mean. This is what the guest lecturer knows. Given all of these facts, (1) is true iff *there is someone in the group who both spoke (1) and is a computer scientist*. This is not the official content of (1); it is not what David said. But it is what the guest lecturer learns; it is the cognitive impact of the utterance on him in this particular situation.

- b) In this case we fix the facts that (1) was uttered by David, that it was in English, that it was of the form “I am an X,” and that the words and construction mean what they do in English. This is what the guest lecturer knows. Given all of that, (1) is true iff *there is a word “K” so that David said “I am a K”, and “K” stands for a certain sort of person, and David is a person of that sort*. This is not the official content of (1); it is not what David said. But given his mumbling, it is what the guest lecturer got from his utterance.
- c) We fix the fact that David is the speaker, and that he is making an assertion. This is what the guest lecturer knows. Given all of that, (1) is true iff *in the language of (1) the words of (1) have a meaning that, given the contextual facts as they are, expresses a true proposition*. This is not the official content of (1); it is not what David said. But given David’s terrible accent, it is what the guest lecturer got from his utterance.

a)–c) illustrate the *relative* concept of truth conditions we need for the epistemology of language. The truth conditions of an utterance u depend on what is *given*. The difference between P and Q above also illustrates this. Q and P provide answers to two different questions. Q answers the question:

Given that an utterance u of “I am a computer scientist” is true iff the speaker is a computer scientist, what are the truth conditions of (1)—what *else* needs to be true for (1) to be true?

P answers the question:

Given that an utterance u of “I am a computer scientist” is true iff the speaker is a computer scientist, *and* that (1) is spoken by David Israel, what are the truth conditions of (1)—what *else* needs to be true for (1) to be true?

Both concepts of truth conditions are perfectly legitimate. And both can be used to get at the states of competent speakers who hear (1). Q gets at what is known of the truth conditions of (1) by those who hear and understand the utterance, but who don’t know who is speaking. P gets at what is known of the truth conditions by those who know all of this and know that the speaker is David Israel.

There are indefinitely many relative truth conditions for an utterance, since there are indefinitely many characteristics C_1, \dots, C_n to plug into the formula “Given that utterance u has characteristics C_1, \dots, C_n , what *else* has to be the case for it to be true?” I’ll use terms of the form “content $_X$ ” to get at some of the more natural sets of characteristics that are useful to think about. I’ll say that the content $_C$ of an utterance is its truth conditions given the facts about its meaning and its context. P is the content $_C$ of (1). Content $_M$ is the content with just the meaning fixed, so Q is the content $_M$ of (1).

Now let’s look at (2). If we fix the meaning of (2), we fix the naming conventions that are being exploited. The content $_M$ of (2) is just P , the proposition that David Israel is a computer scientist. On the account of proper names we have developed, they

are not treated as indexicals. Thus the content_M of (2) is the same as the content_C of (2). So *T* doesn't give us either of the main kinds of content we have identified.

What *T* gives us is the truth conditions of (2) given *some* but not *all* of the facts about the meaning of (2). (2) tells us what else has to be true, given the facts about the meaning of "is a computer scientist" and the fact that "David Israel" is a proper name, but not given the facts about the specific conventions being exploited by the subutterance of this name. We might call this the "content_M of (1) except for 'David Israel'".

This concept of truth conditions, truth conditions with part of the meaning fixed, is very important in the epistemology of language, for it gets at what we know in a familiar and inevitable situation—when we know the meaning of some but not all of the words in an utterance. We use this level of content in planning our utterances, as when we explain the meaning of a word by using it in a sentence. We rely on our interlocutor knowing the meaning of the sentence except for the word. So their grasp of the truth conditions of what we say will be the content_M of our utterance *except for* this word. If we choose our example well, he will be able to figure out the meaning of the statement as a whole. From this, together with his knowledge of what the rest of the statement means, he can figure out what contribution the word must be making, and learn its meaning.

This mechanism is at work in the simple case in which we introduce ourselves. I meet you at a party, extend my hand, and say "I am John". You learn yet another permissive convention for "John", that it can be used to designate me. You start out not knowing this convention. Assuming sincerity, and relying on what you already know, you learn the reflexive proposition, that there is a convention that allows one to call the speaker of the utterance you are hearing "John". You see that I am the speaker, and learn that I am named John. If referentialism is correct, the proposition I express is the metaphysically trivial. But who cares; that proposition is not the one my plan relies on.

On the other hand, perhaps you already know this convention. Suppose you are a new graduate student at the Stanford Philosophy Department. You have seen the name "John Perry" in the catalog, perhaps gone to the library and found an article by me; perhaps you have talked to me on the phone. So you have a notion of a certain person, and realize he can be called "John Perry". When I introduce myself at the party, you assume that the "John Perry" naming convention that I am exploiting is the one you know. What you learn in this case is a fact about the context, that the speaker is John Perry. To get at the cognitive impact of my statement on you—how it differs from my saying "I am I" or "John Perry is John Perry"—we need content_M. You combine this proposition with what you can see, and learn that the person before you is John Perry—the John Perry to whom you have talked on the phone and whose name you have seen in the catalog.

This allows us to see the difference between *indexicality* and *reflexivity*. In the case of non-indexicals, reflexivity disappears when meaning is fixed. Indexicality is the very special case when we have reflexivity at the level of content_M.

Which of these kinds of content, one might ask, is the *real* content of an utterance? The question is misconceived. An utterance has as wide a variety of contents as we

may find useful to isolate, for particular purposes of description and explanation.

We can say that in at least the vast majority of cases, the common sense concept of “what is said”—what we have called “official” content—corresponds to content_C. This is a good reason for an account of content to recognize this concept, but not a good reason to expect it to be the only or even the most theoretically fruitful kind of content.

5 References

References

- [Almog, 1989] Almog, Joseph, John Perry and Howard Wettstein, eds. *Themes From Kaplan*. New York: Oxford University Press, 1989.
- [Barwise and Perry, 1983] Barwise, Jon and John Perry. *Situations and Attitudes*. Cambridge: MIT-Bradford, 1983.
- [Burge, 1973] Burge, Tyler. Reference and Proper Names. *Journal of Philosophy*, 70, 1973, pp. 425-439.
- [Donnellan, 1970] Donnellan, Keith. Proper Names and Identifying Descriptions. *Synthese*, 21 (1970), pp. 335-358.
- [French, et. al., 1979] French, Peter A., Theodore E. Uehuling, Jr., and Howard K. Wettstein, editors, *Contemporary Perspectives in the Philosophy of Language*. Minneapolis: University of Minnesota Press, 1979.
- [Kaplan, 1989a] Kaplan, David. Demonstratives. In [Almog, 1989]: 481-563.
- [Kripke, 1980] Kripke, Saul. *Naming and Necessity*. Cambridge, Mass.: Harvard University Press, 1980.
- [Marcus, 1961] Marcus, Ruth. Modalities and Intensional Languages. *Synthese* 1961 (303-322).
- [Marti, forthcoming] Marti, Genoveva. The Essence of Genuine Reference. *Journal of Philosophical Logic*, forthcoming.
- [Perry, 1993] Perry, John. *The Problem of the Essential Indexical and Other Essays*. New York: Oxford University Press, 1993.

- [Perry, forthcoming] Perry, John. Indexicals and Demonstratives. In Robert Hale and Crispin Wright (eds.) *Companion to the Philosophy of Language*. Oxford: Basil Blackwell, forthcoming.
- [Recanati, 1993] Recanati, Francois, *Direct Reference: From Language to Thought*. Oxford: Blackwells, 1993.
- [Wettstein, 1986] Wettstein, Howard. Has Semantics Rested on a Mistake? *The Journal of Philosophy* LXXXIII, no. 4 (1986): 185-209.