

Lecture One

The Debate Over the Theory of Reference for Proper Names

1. Where we will make our start

This lecture sets the scene for the lectures to come by drawing some morals from, and giving something of a running commentary on, the vigorous debate over the description theory of reference for proper names. It will explain why I dissent from the majority about the anti-description theory message allegedly delivered by Kripke. I don't think he refuted the description theory of reference for proper names. I think he, in effect, told us which version of the description theory we should affirm. Near the end, I will talk about the key distinction, heralded in the prologue, between two quite different things you might mean by the description theory of reference for proper names.

2. The supervenience of reference on nature

I start with something I feel should be relatively non-contentious, although, as we will see, it has contentious implications.

The debate over the description theory of reference for proper names is not about whether the reference of proper names goes by description (in the sense of going by properties, the availability of a *word* or *phrase* for the property that settles reference on a description theory isn't in itself to the point). We know that the reference of proper names goes by description. That follows from the supervenience of reference on nature. If a token name "N" refers to x but not to y , x and y must differ in some way, and in some way over and

2 The Debate Over the Theory of Reference for Proper Names

above the difference that the first but not the second is that to which “N” refers. There must be something about x , a feature of x , which explains why “N” refers to it and not to y .

The point here isn’t one about verification. We aren’t playing a Viennese waltz. True, if *per impossible* “N” referred to x and not y , in the case where x and y are exactly alike, one could never find out which of the two “N” named. But what is crucial is not the epistemological point but, rather, that there would be nothing to make the difference, and a difference in terms of what is referred to isn’t arbitrary. Insisting that there might still be a fact of the matter as to which “N” named would be like insisting that exactly one of:

Had Bizet and Verdi been compatriots, they would both have been French

and

Had Bizet and Verdi been compatriots, they would both have been Italian

must be true. The mistake in insisting that exactly one of the above sentences must be true, no matter how symmetrical the situation may be, isn’t to do with verification problems (though they would exist, of course); it is to do with arbitrariness.¹

The difference in nature between x and y when “N” refers to one and not the other may be in their relational properties. Tweedledum and Tweedledee are, we may suppose, qualitative duplicates (not that they are in Lewis Carroll, *Through the Looking Glass*; they are same sayers and maybe mirror images). All the same, they have different names. “Tweedledum” will refer to one of them and not the other; ditto for “Tweedledee.” The explanation will lie in the different relational properties they possess. To say exactly which differences are crucial would trespass on matters for later discussion. But among the differences are the following: Tweedledum but not Tweedledee was baptized “Tweedledum”;² Tweedledum but not Tweedledee

¹ See Goodman (1947).

² I mean “baptize” in the sense of giving a name to, not in the sense of admission to the Christian Church. See Kripke (1980).

answers to the name “Tweedledum”; Tweedledum but not Tweedledee is the object that sentences containing “Tweedledum” carry information about; and so on. The fact that “Tweedledum” refers to Tweedledum and not to Tweedledee is not a counter-example to the supervenience of reference on nature but instead illustrates the fact that the nature in question has to be understood so as to include the relations between what is referred to and the token name in question (excluding of course the very relation of being what the name names – that would trivialize the supervenience of reference on nature). We might also have a case of qualitative identity with difference in reference where the name is one and the same. Paris Texas and Paris France are very different. But they might have been qualitative duplicates. Consistent with this, it might have been the case that some tokens of “Paris” refer to the first city, and other tokens of “Paris” to the second city. There will, in that case, be a difference in the relationship between the respective *tokens* and the respective cities. Causal theorists will tell us that the difference lies in which Paris is in a certain causal connection to which tokens. We will be saying something similar in due course, but from a perspective friendly to the description theory.

It follows from the supervenience of reference on nature that, whenever a name refers to an object, there is a property of that object such that the name refers to an object if and only if the object has the property (in the inclusive sense of a property that includes relational properties). In that sense, a description theory of reference for names has to be true.³ What then is the debate about?

3. The availability issue

The debate is – can only be – about the *availability* of the property or description that settles the reference. Supervenience is compatible with unknowability. If names were assigned to objects by some secret process constrained only by a respect for supervenience – imagine that names are paired one-one with uniquely possessed properties by an evil demon who won’t talk, and that each name refers to the object with the property the demon pairs the name with – then the

³ The same goes, of course, for token indexicals and demonstratives.

reference-determining property for any given name would be unavailable (except to the demon, and he's not talking), but supervenience would be respected. There would no difference in reference without a difference in referent.

But the anti-description theorists' position isn't – had better not be – that the relevant property is not available to anyone ever. Unless we are to give up entirely on finding a theory of reference for proper names, we are committed to the availability, to some person at some time, of a truth of the form: “*N*” refers to *x* if and only if ... , where the ellipsis is filled with a property specification or description that says how *x* must be to be that to which “*N*” refers. That's what a theory of reference for a proper name *is*. For example, those who say something like, “‘*N*’ refers to *x* if and only if *x* was baptized ‘*N*’ (or with a name with so and so a relation to ‘*N*’) in the past, and speakers of the language have agreed to use later tokens of ‘*N*’ in such and such a way” are giving one account of how to fill in the ellipsis; those who say something like, “‘*N*’ refers to *x* if and only if ‘*N*’ is associated with a property *P* that *x* alone has” are giving another account of how to fill in the ellipsis. But the framework is the same in both cases. To provide a theory of reference *is* to provide the property that determines what a name refers to; it is, that is, to provide the conditions under which a name refers to what it does refer to.⁴ This means that to hold that we will never know the property that determines the reference of a name is to hold that we will never know the correct theory of reference for names. We might as well stop searching right now.

It follows that the only plausible version of anti-descriptivism affirms that the reference-determining property for a name is *relatively* inaccessible. The version allows that some know the property, namely, those in possession of the correct theory of reference for the name, but many, perhaps most, don't know the reference-determining property. Thus many of those opponents of the description theory who favor some version of the causal theory respond, when asked why their theory isn't a kind of description theory – the kind known as causal descriptivism⁵ – that the causal story is not known by the folk. We know it, they say, but most do not, and give

⁴ Cf. Robert Nozick's point mentioned by Kripke (1980, p. 88, *n.88*.)

⁵ For causal descriptivism see, e.g., Kroon (1987; 2009).

as examples the names “Feynman” and “Jonah.” They argue that certain tokens of “Feynman” in the mouths and from the pens of the folk refer to Feynman, the physicist, but the folk do not know any property unique to Feynman, and hence do not know the reference-determining property, for that property must be unique to Feynman.⁶

The trouble for this “we know but they don’t” position is that the folk *do* know a property unique to Feynman. The folk know perfectly well that sentences like “Feynman worked at CalTech” carry information about Feynman. This is obvious from the use they make of these sentences. This knowledge is as commonplace as the knowledge that what appears on television screens carries information (or putative information, but I will often drop the “putative”; it gets tedious). What is more, they know that information transmission is underpinned by causal connections. Someone watching a tennis match on television does not need to be told that the information they are getting about the match depends on causal links between what is happening on court and what is happening on the screen. What is more, the folk know that the information about an object carried by sentences containing proper names typically depends on causal chains initiated by the assignment of a name to the object. They know that the information carried by names requires *naming*. What is more, they know the information carried relies on their language community adopting conventions of usage that mean that sentences containing names preserve information about what is named.

What makes it especially clear that the folk know all this is what they *do* when they hear a sentence like “Feynman worked at CalTech.” They repeat it and, being good citizens, they would hardly repeat it unless they thought it carried information, and they would be slightly bemused if asked who or what it carried information about – “The person called ‘Feynman’ of course.” If the folk happen to know that CalTech is in Pasadena, they may well go on to utter the sentence “Feynman worked in Pasadena.” In doing this, they are participating in the conventions that underlie the transmission of information using names, and, when they do it, they don’t need a tutorial from philosophers of language first.

This is very much in line with what Kripke says (see especially 1980, p. 91). Although he doesn’t say it this way, Kripke’s account

⁶ See, e.g., Devitt (forthcoming); Devitt and Sterelny (1999); Soames (2005b, p. 299)

is one that makes proper names part of what Timothy Williamson calls “a channel for the acquisition of knowledge” (2007, p. 264).⁷ And it is this picture that tells us why “Madagascar” refers to Madagascar and not the part of Africa originally named “Madagascar.” Sentences of the form “So and so happened in Madagascar” carry information about Madagascar, not about the part of Africa. (I assume that Evans’s (1973) account of the history of the use of “Madagascar” is correct; if it isn’t, please pretend that it is.) Had there been a major earthquake in Madagascar last night, today’s papers and internet sites would contain sentences like “Many died in Madagascar overnight” and “Relief pours into Madagascar.” However, sentences like these would not have appeared had the major earthquake taken place in the relevant part of Africa. None of this is news to the folk – surely. Or suppose I hear the name “Napoleon” and decide it would be a nice name for my pet aardvark, I don’t refer to Napoleon even if the token I heard was a use that referred to Napoleon.⁸ The reason is that my uses of “Napoleon,” stemming from the decision to use it for my pet aardvark, will carry information about the aardvark and not Napoleon.

A way to bring this out is to imagine that I am quite wrong. Imagine, that is, that a causal, information-preserving account of the reference of proper names that appeals to the role of naming, plus observing the convention of transmitting information using names, is correct but is not known by the folk. Information is valuable. People want it and pay for it. If names are key sources of information about objects but this is something known only to certain philosophers, we, or they, should be shouting the news from the roof tops, or, at the least, sending e-mails to colleagues in other departments passing on the good news. How could we reasonably restrict ourselves to talking about it in philosophy seminars? But of course no philosopher is going to send an e-mail to, say, the Politics Department in their university telling them that the appearance in the papers of sentences like “Gordon Brown increased funding for Barclays Bank yesterday” carries information about Gordon Brown, and does so in virtue of an information-preserving causal chain involving the names

⁷ See also Kroon (2009), Evans (1973; 1982, ch. 11); the basic idea can be found in many places, as you would expect if it is a bit of folk wisdom.

⁸ The example is a modification of one of Kripke’s (1980, p. 96).

given to Brown and to Barclays, and the observing of conventions for the transmission of information. (Imagine the reaction of the Politics Department.)

Similar points apply to the famous “Jonah” example (Kripke 1980, p. 67f.). Many causal theorists distinguish themselves from causal descriptivists by insisting that, although the reference of certain tokens of “Jonah” go by a certain causal property related to a certain naming some time in the past, this property is only known to a few. In particular, it isn’t known by the folk. However, as one would expect, there has been considerable interest over the years in whether Jonah, the Jonah of the *Bible*, existed, and if he did who he was. The same goes for Helen of Troy. Scholars pursuing these questions have never felt the need to employ philosophers of language to help them in their search for the Jonah of the *Bible* or for Helen of Troy. But the task of finding who, if anyone, Jonah or Helen were is nothing other than the task of finding who, if anyone, fits the bill to be the referent of “Jonah” or “Helen” in certain texts, and fitting the bill is having the right properties. It seems the folk do know the properties that determine the reference of names like “Jonah” and “Helen of Troy” in certain texts. What is more, we can all give a rough characterization of what scholars seeking Jonah or Helen do. They start with the token names in the texts in question and work backwards, looking for the causal origins of those tokens in the story-telling practices of the past. When they find the right kind of causal origin, they have found their man or woman. That’s the methodology of books and television series devoted to asking who, if anyone, was Helen of Troy.⁹ But – to labor the point – discussions of who she was aren’t especially the preserve of philosophers. Or, to give an example familiar to many parents, they get asked by their children, “Did Robin Hood exist?” and, “If he existed, who was he?” They know that the answer to these questions lies in the causal origins of tokens of the name “Robin Hood,” in various texts, films, and television series, and the relations of those origins to information conveyed by sentences using the name.

Our final example to bring out the point – and apologies for the mouthful that is coming – that this *information channel based on*

⁹ I have heard it said that it is more likely that Jonah existed than that Helen of Troy did, and of course it may turn out that there is no determinate answer, in one or both cases.

conventions for transmitting information using names account of the property that determines reference is folk theory is the way sentences containing proper names assist us in known ways to navigate our world, and the effects the names have on where we end up. A person's coming across a sentence containing a name can dramatically change the probabilities of that person being at, or not at, the thing named. Someone's hearing "Sammy's is a good restaurant" increases the chance that they will be at Sammy's. Someone's reading "Chernobyl is radioactive" lowers dramatically the probability of their going to Chernobyl. Or consider a phenomenon especially familiar to academics. Reading an e-mail containing the sentence "There is an interesting conference coming up in Bielefeld" increases the probability of faculty from all over the world ending up in the one place, namely Bielefeld (an early version of some of these ideas was presented at a conference in Bielefeld).

What is the explanation? Part of the explanation has to be that we know of distinctive marks of the places in question. You keep clear of Chernobyl by knowing a property, a marker if you like, distinctive of Chernobyl, and making sure that wherever you are doesn't have that property. You get yourself to Bielefeld by knowing a marker for that city, and making sure that where you end up has that marker. Of course each of us doesn't need to use the same marker, but we each need a marker. We each need a uniquely possessed property but don't need the same one. Why do we need a *uniquely* possessed property? Suppose that the best I can do by way of pinning down Bielefeld is with a property that it shares with Berlin. Then it will be 50:50 or thereabouts whether I turn up in Bielefeld or in Berlin. But we all know that it isn't 50:50. Everyone who went to the conference in Bielefeld knew in advance that they would very likely get there, and knew that the same held for the others who set out for Bielefeld from all over the world, barring accidents.

How come we were all so confident? Certainly, many of us knew very little about Bielefeld when we first read the e-mail and were well aware, for instance, that there might be two or more cities called "Bielefeld." What we did know, however, was how to use the word "Bielefeld" to get the information we needed. We sent e-mails containing the word in sentences like "How do we get to Bielefeld?" to the organizers who sent out the original e-mail. We used the word in talking to travel agents and colleagues. And we knew that the

sentences that came back to us would carry information about Bielefeld. That is to say, we exploited our knowledge that the city was the object about which these sentences carried information. And this was all common knowledge. None of us felt it necessary to send out e-mails explaining about naming practices and the way they sustain information channels in order to justify using the word “Bielefeld” to get the information we needed. True, many of us were philosophers of language, but historians and bankers are equally good at using names to get to the city where the conference they want to go to is being held.

Speaking more generally, it is common knowledge that the world is a complex place with a huge number of different objects. We often need to be able to differentiate one from another – I want my mail to get to my house and not yours – and although it is plausible that there will be some property unique to each and every object – there are no *absolutely* identical twins, for example – this fact is of little use to us if we don’t know the differentiating property. Of necessity we set up systems that differentiate in known ways between objects. That’s what we are doing when we name streets, cities, and people, give numbers to rooms, put the black dot on one of the two white cue balls in billiards, and so on. Often (tokens of) the names and numbers – the differentiating labels – are physically attached to the objects. For example, room numbers are typically placed on or very near the room they serve to differentiate. But there are plenty of exceptions, our relation to our own names being an obvious example. In hospital, where it is especially important that there be no confusion, we are tagged with our names – they are literally attached to us – but by and large our connection to our names is a causal one. I *answer* to my name; I *put* my name on forms. But assigning names is only a first step in differentiating objects. There are many cities called “London.” How then do we ensure that a letter gets to the desired city out of, as it might be, London England and London Ontario? One way is to append the word “England” or the word “Ontario” to “London.” But the most widespread way of handling the problem is to trade on our knowledge of the way people observe the convention of transmitting information using names via their causal connections to the named object. A meeting is called for room 110 next Friday. Where do I go? There are many rooms numbered 110 in my university, and maybe millions numbered 110 around the

world. But I know that the person calling the meeting is observing the convention of using “110” to deliver information about the room in question. This in itself may answer my question. I know there is only one room 110 that stands in the kind of causal connection that would allow her to observe the convention. But if this isn’t enough, I send her sentences containing “110” and get the detail I need from her responses, for I know that she will be observing the convention, and thus know that what will come back will likely be identifying information about the room I need to find.

What I have said in the last few paragraphs is common knowledge, folk knowledge. It is part of understanding the role of names in a natural language like ours. The property that settles the reference of a proper name isn’t an unknown one. It is the kind of property Kripke talks about. My dissent from the causal theory of reference is simply my insistence that the property is no secret.

I have said things like this before; others have said things like this before.¹⁰ We have had less impact than one (I) might have hoped. It is still conventional wisdom that the description theory of reference for proper names is false.¹¹ One explanation for my failure is no doubt that I did not make my points well enough. But there are four other explanations for resistance, each of more interest.

One is that many insist that the question of folk knowledge of the reference-determining property is incidental to the central question. Maybe the folk do know the property and maybe it is the property that causal theorists of reference say it is, but the key point is that a user of a name can refer to the referent of the name without knowing the property. Their claim is that user ignorance is compatible with successful reference.¹² Another reason for resistance to our defense of descriptivism is that many hold that I am wrong in insisting that the folk have knowledge of the reference-determining property. Sometimes the argument is that the folk, and highly trained philosophers if it comes to that, cannot *specify* the reference-determining property in full, exceptionless detail. What the folk have is a recognitional ability subserved by information at the sub-personal level.¹³

¹⁰ See, e.g., the papers by Kroon cited above; Evans (1982); Lewis (1999); Jackson (2005).

¹¹ Thus the sentence starting “But with the demise of descriptivist theories of proper names, ...” in Tye (2009, p. 81).

¹² See, e.g., Devitt and Sterelny (1999), and Soames (2002; 2005)

¹³ See, e.g., Davies (2004) and the response by Braddon-Mitchell (2004).

They don't have personal level knowledge of the reference-determining property. A third resistance point is that some insist that a theory of reference for names must make sense of the way quite young children refer using names, and it is hard to believe that they know about information channels sustained by naming and usage conventions. Finally (and most influentially is my impression), many hold, first, that the description theory of reference is committed to the equivalence, or the interchangeability in a range of contexts, of sentences of the form "*N* is *F*" with sentences of the form "The *D* is *F*," or maybe the "The actual *D* is *F*," or maybe "The *D* is *F*," combined with a special scope rule, where "*D*" gives the reference-determining property; and, second, that there are decisive objections to equivalence and interchangeability.¹⁴ They are right on the second point. There *are* decisive objections to equivalence and interchangeability. But it is a mistake to hold that the description theory is committed to them, a mistake sometimes made by description theorists, it should be said. (*Mea culpa* but I wasn't alone.) We will discuss the issues raised by these four explanations for resistance in the order I have listed them.

4. What is required of a user of a name?

The objection, remember, is that though I might (*might*) be right about the reference-determining property and folk knowledge of it, a user's use (call him "Fred") of "Cicero" can refer to Cicero when Fred doesn't know the property.¹⁵ How so? Is the idea that a certain, information-bearing causal chain from Cicero to "Cicero" in Fred's language community is, in itself, enough to make it the case that Fred's use of "Cicero" refers to Cicero, regardless of Fred's mental state? But now there is a crucial question to be asked: is the token use of "Cicero" *by Fred* part of that causal chain or not? Suppose not. Then there is a serious problem. There are very many causal chains from very many Ciceros to very many tokens of "Cicero."¹⁶

¹⁴ See, e.g., Soames 2002; 2005a; 2005b.

¹⁵ For versions of this line of objection see Devitt and Sterelny (1999, p. 99) and Soames (2005, p. 301).

¹⁶ I take this point from Kroon (2009, p. 152), but please do not hold him responsible for my development of it.

Which Cicero does Fred's use of "Cicero" name? The defender of reference in the face of ignorance has no way to answer this question. Any answer will be as arbitrary as any other answer. The only way out is to include Fred's use in the causal chain. That way there will arguably be a non-arbitrary answer to the question of what his use names. It will name the same thing as its predecessors in the chain that Fred's use is a part of – say, one that goes back to Cicero, NY. But now we come up against Kripke's aardvark point. The use has to be the right use. If Fred hears the word "Cicero" used by someone to refer to Cicero, NY, and thinks it would be a good name for his aardvark, Fred isn't referring to Cicero, NY, by his consequent uses of "Cicero" despite its causal connection to a token that does refer to Cicero, NY.

What is the right kind of use? Surely, the use where Fred is following the convention of conveying information about something named in the past. The use on which, if he hears "Cicero is a fast growing city," he will repeat it, or maybe employ his knowledge of the relation between house prices and growth, to pass on putative information with the sentence "House prices in Cicero may well have held up better than those in Florida." The attempt to bracket out Fred's mind from the story about what he refers to using a name fails. It has to fail. The point was in fact clear from the moment we came across the aardvark.

It might be suggested that, although we can't leave Fred's mind out of the account of how "Cicero" in his mouth names what, if anything, it names, we do not have to suppose that he knows the linguistic conventions governing naming (which is in effect what we required above). Much less is enough. He needs merely to have the right intentions. The suggestion might be that Fred may succeed in referring to the very thing others in his community refer to when they use the name simply as a result of the name's being in public use, together with the standing intention that his use of words conform with the linguistic conventions of his community.¹⁷ Here what must be meant is the standing intention to use words in conformity with some given *token* use or uses. Otherwise we will have the "too many referents" problem over again. There will, that is, be no answer as

¹⁷ Something like this suggestion seems to be at work in Soames (2005b, p. 301, and 2007, p. 39).

to which of the many Ciceros, Fred's use of "Cicero" refers to. Let us suppose this is what Soames has in mind, and suppose further that the token use Fred has the standing intention to conform to is one that refers to Cicero, NY. There are three cases. In the first, he has no idea how his community uses the token of "Cicero" in question. In this case he won't be using the word very much. He will be like someone with the standing intention to observe the conventions of a bridge bidding system in the absence of knowing what they are. People in that situation don't do any bidding until they learn what the conventions are. The very fact they have the intention to conform means they don't act until they know, or at least have an opinion about, what it is to conform. *Mutatis mutandis* for those who seek to have their word usage conform to that of their community. In the second case, Fred is confident that the token of "Cicero" in question is used for a certain color (as some tokens of "Cicero" are, I understand). He is wrong, but that is what he thinks. His standing intention then makes him use "Cicero is attractive" just when he thinks that the color is attractive, and "Cicero is becoming more popular" just when he thinks that the color is becoming more popular. It is obvious that Fred is not referring to Cicero, NY. He is instead referring to a color. His standing intention in this case fails to bring his reference into conformity with that of his language community. In the final case, he is confident that the token in question is used as part of an information preserving chain, etc. In this case his standing intention will mean that his use of "Cicero" refers to Cicero, NY, but this will be *because he knows about the convention for using names and is following it*.

What is the moral? Intending to conform and having standing intentions to conform aren't, in the final analysis, what matter. *Conforming* is what matters. And in a way that was obvious from the start. Why seek to conform or have a standing intention to conform if *actually* conforming isn't, at the end of the day, what is crucial?

We should, however, note that there is a sense in which the emphasis the suggestion places on the role of one's language community is fully in accord with the kind of description theory I affirm on behalf of the folk. It is users' language communities, with their observance of the conventions, that underpin the informational integrity of the causal chains that run from dubbings of objects to token sentences.

Our difference is that I insist that users of the name know about the conventions and the information-preserving chains, and need to know about them, if they are to refer when they use a name. Here it is apposite to recall a passage in *Naming and Necessity*:

In general our reference depends not just on what we think ourselves, but on other people in the community, the history of how the name reached one, and things like that. It is by following such a history that one gets to the reference. (Kripke 1980, p. 95)

My (and I am sure I speak for other causal descriptivists here) objection to this passage is restricted to Kripke's suggestion that depending on "what we think ourselves" would leave out the role of the community. The role of the community is common knowledge. It is something we think ourselves. This is clear, as we said earlier, from the way the folk follow "such a history" in seeking Helen of Troy or Robin Hood.

5. The issue about personal level knowledge

Our brains locate sounds by using the connection between the location of a sound and the out-of-phase effects at our ears when the sound reaches them (and other stuff like the differential filtering effect of the head, but we will simplify). But we don't draw conclusions about the location of a sound by knowing the out-of-phase effects at our ears and proceeding to infer the location from this knowledge. The processing is sub-personal. Something similar is true for face recognition. We rarely identify someone by noting that their face has property *P* and inferring that, as only George's face has property *P*, it must be George we are looking at; we simply recognize George's face. The recognition is undoubtedly driven by a property distinctive of his face, but the property is recorded at the sub-personal level.

There are exceptions of course. *Descriptions* of wanted persons sometimes lead to their capture. Perhaps the bulletin includes the words "the suspect has a tattoo of a rose on his left cheek and a fresh scar on his forehead," and that is enough to enable members of the public to identify the suspect. This would be a case of using personal

level information to identify someone, but this is not normally how we identify our friends, or politicians on television. Try giving a list of properties of Tony Blair's face that enables someone else to identify him on television with confidence. You very likely won't be able to, but you can recognize Blair easily enough.

Is the property that determines reference like the property that underpins face recognition, something that operates at the sub-personal level? If that is right, the folk don't know the property. Their brains 'know' it if you like, but they don't. However, this suggestion can't be right.

First, if it were right, *philosophers* aren't going to find the reference-determining property. It took some physics and neuroscience to find the property that underpins our ability to locate sounds. Philosophical reflection didn't do it, and no one thought for a moment that it would. Now, we know that there is a reference-determining property – that follows from supervenience – and some philosophers claim to know what it is. That is precisely what the causal theory of reference for proper names is all about, and of course the view I am defending is a version of the causal theory with an emphasis on information-preserving channels. But if which property it is is a sub-personal matter, we should be calling in neuroscientists and cognitive psychologists. Philosophers *qua* philosophers cannot tell us, any more than the folk can. It would be a mistake in principle to think that the seminal philosophical ideas we find in *Naming and Necessity* bear on the question. This is very hard to believe.

Secondly, when we recognize in ways underpinned sub-personally, we typically need to be 'trained up' beforehand. The first time you see Tony Blair on television you won't recognize him. You need to be told that it is Tony Blair you are looking at it. You need to get the information at the personal level. *Subsequently*, you are able to recognize him without personal level cues (though they often help). Also, we are all familiar with the way our recognitional capacities are blunted by unfamiliar surroundings. Someone we recognize on the tennis court we may fail to recognize in the police station. However, our ability to find cities using sentences containing their names does not depend on being acquainted with the cities first (no 'training up' is required), and is surprisingly resilient. Indeed, one of the great strengths of our naming practices is the way they enable us to find *unfamiliar* places in *unfamiliar* situations. We find Bielefeld

or 10 Elm Street despite never having been there before, and it doesn't matter much what the weather is like.

Thirdly, the key fact that tells us that an identification underpinned by the sub-personal is in fact underpinned by the sub-personal, namely, our inability to make the identification from descriptions alone, is missing in the case of finding the reference of proper names. Readers and reviewers of books or television series devoted to arguing about the identity of Helen of Troy or Robin Hood are in a position to judge whether or not the book or program succeeded in identifying Helen of Troy or Robin Hood, and, if they are undecided, this will be because either they need more information – more personal level information, that is – or because the matter is indeterminate: there is no clearly correct candidate to be the right kind of causal origin of the information carried by “Helen of Troy” or “Robin Hood” sentences in the classic texts. But the evidence offered by the books and television series is essentially in words on the page or screen, or from the mouths of those fronting the cameras. And of course I could have made essentially the same point by reference to the very literature that drives so much of the debate over the theory of reference, starting with Kripke's and Putnam's seminal contributions. The literature is full of *descriptions* of possible cases combined with invitations to make one or another judgment about what thing is referred to by some name or other. This methodology would be a nonsense if our judgments were underpinned by the sub-personal.

Finally, the importance of, for example, ending up in the right city for a conference means that it makes very good sense that we should have personal level markers for cities, and the same goes for people and the other objects we name. The rationale here is the same as that for having personal level markers for the poisons in our medicine cabinets. We make sure our children know what the skull and crossbones on a label signifies. (We return to this issue in the second half of the final lecture.)

6. The demand for precise and explicit specifications

I said that the folk know the reference-determining property. I gave a sketch of what it was and noted its importance for the acquisition of information about objects. I supported the claim that the folk

know the property by observing that, if they don't, it is strange we aren't telling them about it more vigorously given the importance of information, and by observing, moreover, that the way folk use sentences containing proper names makes it clear that they do know the property. However, I did not give a precise specification in words of the property. Some insist, or seem to insist, that any defender of the description theory must do what I did not do.¹⁸ We mentioned Devitt and Soames in this connection. (It is perhaps worth noting that some who hold the view discussed a moment ago, the view that we don't have knowledge of the reference-determining property but do carry, at the sub-personal level, information about it, may be assuming that all personal level knowledge can be given a precise verbal expression. But this would be a mistake. Dogs and gorillas have beliefs but cannot give them precise verbal expressions. All the same, belief is a personal level state.)

A short reply on behalf of the description theory is that knowing a property and being able to give a precise specification of the property in words are two different things. But, as the issue comes up in discussion so often, the objection deserves a longer discussion. However, this longer discussion needs to be set in the context of some remarks on what a theory of reference for names is a theory *of*.

7. What is a theory of reference a theory of?

To date, we have followed the common practice of taking for granted our target notion – the reference relation between names and the things they name – and have moved straight to a discussion of the way the relation supervenes on the nature of what is named or referred to, and of the arguments against the account of the relation given by the description theory. But what, more exactly, is our target? Obviously we are talking about a word to world relation but there are many word to world relations. Which word to world relation is the one we are focused on, which is the one that has been the focus of so much discussion since Putnam's and Kripke's contributions?

¹⁸ Sometimes the objection is put by saying that description theorists cannot *supply* the description; see, e.g., Devitt (forthcoming). But of course they can: "the reference-determining property for 'Plato'" is a perfectly good description.

That is to say, what is a theory of reference a theory *of*? (Thanks here to Josh Sheptow for discussion.)

Let me start with a simple distinction. I don't know the day I will die. Does this mean I don't know the reference of "the day I will die"? Yes and no. The sense in which I *don't* know the reference is that I don't know which day is the fateful day. I know it will be after today. I know it will be before 2040. But if you listed the days between now and January 1, 2040 and asked me which one it is, I couldn't say (although I do not give much credence to the days near the end of that period). The sense in which I *do* know the reference is that I know what it takes to be the reference of "the day I will die." I know how a day has to be, the property it has to have, to be the day I will die. In that sense I do know the reference of "the day I will die."

Our interest in these lectures is primarily with reference in the second sense. We are concerned with the way we use words and sentences to represent things as being a certain way, or, equivalently, to give putative information about how the world is. And what it takes for x to be the referent of "the day I will die" is what we represent about how things are, the putative information we give about how things are, when we affirm " x is the day I will die." Or consider the sentence "Today is not the day I will die." If one knows what it takes to be the reference of "today" and what it takes to be the reference of "the day I will die" (and the way negation and predication work), one knows what that sentence represents about how things are. It says that the thing which has what it takes to be the referent of "today" lacks what it takes to be the referent of "the day I will die." Our interest is in the content of language in the sense tied to how sentences represent things to be, the putative information they deliver, with, as we will often abbreviate it, the *ir-content* of sentences.¹⁹ This means that our interest in the reference of words lies in how something has to be to be the reference of a word, for it is that which feeds most directly into how sentences containing the words represent things to be, into, that is to say, their *ir-content*. More especially, we are interested in the contributions of names of kinds and proper names to the *ir-content* of a small number

¹⁹ Sometimes I will simply talk of content when the context makes clear that it is *ir-content* that we are discussing.

of rather simple sentences containing them – sentences like “Plato was a great philosopher,” “The conference is in Bielefeld,” and “Water is nearby.”

In this sense of reference, the question, What is the reference of “circle”? is a boring question. Boring because we all know the answer. We all know what we are saying about how things are when we use the word “circle” to say how something is. We are saying that it is an instance of the following shape: ○. To say it in terms of reference: there is a function, the reference function, that goes from “circle” to the property of being a circle (or maybe the set of circular things in logical space, or ... ; the different ways of putting the same basic idea are not significant in this context), and in understanding the word “circle,” we grasp that function. In this sense of “reference,” there aren’t many reference relations between “circle” and the world. There is just the one, the one from the word to the shape. Any other word to world relation delivers the wrong answer to be reference in the intended sense. Or better, delivers the wrong answer as far as English is concerned. There is no reason why there shouldn’t be a language that uses “circle” to do the job we use “square” for, and *for that language* the reference relation we are talking about goes from “circle” to being square rather than to being round.

We will sometimes talk of the ir-content of a word, “circle,” as it might be – the word, not a sentence containing the word – when it is important to bear in mind that we are talking of the reference of the word “circle” in the sense that relates most directly to its contribution to the ir-content of, say, “There are circles.” In these terms, the boring question is, What is the ir-content of “circle”? However, a question that is far from boring is, How come “circle” has that content and not, say, the content that the word “square” has? Even for a simple word like “circle,” the *getting* of content (reference) question, as we might call it, is challenging and controversial. The same goes *mutatis mutandis* for sentences containing the word. Saying that the ir-content of “There are circles” is that there are things with the shape – ○ – is unexciting, because it is so obviously true. But how that sentence comes to have the ir-content it does have is highly controversial.

It is perhaps worth taking a moment to point out that there is no significance in the fact that I have chosen to frame matters using “reference” instead of, say, “satisfaction.” The key point – a point

we will discuss at more length in the next lecture – is that our ability to use language to say how our world is shows that there are word to world relations we grasp, just as our ability to use the map of the London Underground to find our way across London shows that there are map to distribution of stations relations that we grasp. I have expressed what we grasp as the reference function from, e.g., “circle” to shape $\bigcirc$, but I could equally have said that what we grasp is that an object *satisfies* “circle” just if it has the property: shape $\bigcirc$. And of course there are still further ways one might express the basic idea.

I can now say what it is we are discussing when we discuss the theory of reference for names. We are discussing the role of names in making claims about how things are. I think Jason Stanley is right when he says “Suppose that Hannah utters ... ‘Bill Clinton lived in Arkansas’ ... Hannah imparts certain information about the world” (2007, p. 5). To say it in the jargon of representation, by uttering “Bill Clinton lived in Arkansas,” Hannah represents, correctly as it happens, that things are a certain way; that’s the (putative) information she imparts. Mostly, in the interest of keeping things as simple as possible, we will focus on sentences of a simple subject-predicate form. So for us the question, What is a theory of reference for proper name “N”? is in effect the question, What is the contribution of “N” to how “N is F” represents things to be, to the putative information the sentence delivers?

From this perspective, the debate over the right theory of reference for names is the debate over the contribution that names make to how things are being represented to be. It is over their contribution to the informational value of sentences containing them. The version of the description theory I have been defending says that their contribution is to be parts of information channels. By way of contrast, one version of the direct reference theory would say that the contribution of names is to pick out their referent and *that’s it*. This kind of direct reference theory typically allows that there is such a thing as the reference-determining property, and it may or may not allow that the folk know what it is, and the property may or may not be as a version of the causal theory of reference says it is, but what the theory holds firm to is that the property in question is not part of the content. Instead it is the content determiner. In terms of the distinction alluded to a little while ago, the answer to the question, What

is the content of “N”? is that it is some object *o*, whereas the answer to the question, How come “N” has content *o*? is in terms of the property.²⁰ One version of this kind of direct reference theory says that “Plato” has the content Plato, and the reason why it has that content is that Plato and “Plato” stand in a certain causal, information-bearing relationship.

We can now see what is wrong with the objection that supporters of the description theory have not specified the relevant property with due precision. One thing is clear. We *are* able to describe the world using words. We can say what things are *like* using words. Why am I so upset when I hear the word “terminal” in connection with a disease I have? Because I know the property being ascribed (for I know what it takes for “terminal” to apply to something), and it is one I don’t welcome. Why were Democrats so pleased to hear the words “Obama has won” uttered on television on the night of November 4, 2008? Because they knew the property being ascribed, and it was one they welcomed. Why do we bother saying things like “Sweden is a liberal society”? Because we know what we are saying about Sweden and it is something worth saying. Our ability to describe the world using words is a datum.

This means that for some (*some*) words it is a datum that the description theory is true. There are words that we know apply to *x* just if *x* has a given property.²¹ However, it is not possible to give a precise expression in words of what it takes to be a liberal society. We can say some useful things, and often do so when we explain to our children what makes for a liberal society. Having a free press is important, as is having a right to vote that operates in a way that gives more than one party some chance of winning, whereas a

²⁰ Stalnaker has been as clear as anyone I know in putting this kind of position, see, e.g., (2003a).

²¹ I know philosophers of language sometimes seem to be denying this. They produce sentences like “The reference of a word ‘W’ is not settled by known properties P_1 , P_2 , etc. Whether ‘W’ refers to *x* is always a matter involving best theory and maybe selectional history,” in a way that might suggest this is a thesis for words across the board. But this would not be the right way to read them, for two reasons. First, they use the words “ P_1 ” and “ P_2 ” in a way that presupposes that we do know the reference of those words. In expressing the view that certain properties don’t settle reference, they perforce use *words* for precisely those very properties. Secondly, there would be no point in using terms like “best theory” and “selectional history” unless they and their readers knew which properties had to be instantiated in order for those terms to apply.

country's being cold in the winter is in itself neither here nor there. What is more, we can say that someone who looks for the property of being a liberal society as if it were something separate from having a free press, equality before the law, and so on is confused. Being a liberal society is a matter of some suitable combination of equality before the law, the right to vote, and so on. It isn't an extra feature that somehow, mysteriously 'goes along with' equality before the law, the right to vote, and so on. However, it is notorious that we cannot write down a precise list of necessary and sufficient conditions for being a liberal society. Opinions differ as to why this is impossible and about its significance, but there is consensus that it is impossible.²²

The same point can be made with terms for shapes. For some shapes we know the formula. The word "circle" applies to a shape just if it is a closed plane figure, each point of which is equidistant from a fixed point. In that sense we can give the description or property that determines its reference. However, there are shape terms for which this isn't possible. Those who can read print and handwriting in the Roman alphabet have at their disposal the term "an a." It applies to the relevant commonality between: a, a, A, A, ... Its reference goes by description: it applies to something just if it is of relevant qualitative kind. We know what the kind is. This isn't a case of sub-personal knowledge. When someone scrutinizes a sample of handwriting and declares a certain configuration to be an a, you know at the personal level – you know, that is, in the ordinary sense of "know" – what they are saying about how the configuration is. But we cannot give a neat set of rules for when a shape is an a. This is one reason it is hard to program computers to read handwriting (and printouts in non-standard fonts). Or think of the knowledge one gets from perception. Perception often yields knowledge about the distribution of *properties*, but our ability to put that knowledge into words is limited.

The moral is that we should not demand that supporters of a description theory of reference for a word, "W," give what is in effect a reductive analysis of the property they hold determines the reference of "W." They must be able to give a reasonable indica-

²² For one opinion about why it is impossible and its relative insignificance, see Jackson (forthcoming a).

tion of the property in words – as we did for being a liberal society, and for being part of an information channel – and it is important that we can recognize cases by the descriptions given of them, but that is possible for the reference of both “is a liberal society” and “Gödel.”

I guess, to put it in slogan form, the point I’m making is, beware the old double standard.

8. In the mouths of the very young

Very young children parrot our words. Sometimes they do more than parrot them; they repeat them with some degree or other of understanding. All the same, it would be strange to make their use of a word a major constraint on a theory about *our* use of it. It would be like insisting that what an expert means by “RAM” has to conform to what I mean by it. But something awfully like this does seem to be going on in the objection to the description theory of reference that runs somewhat as follows: very young children refer to Paris when they use the word “Paris” as in, say, “Mummy and Daddy are in Paris,” but it is obviously wrong to credit them with a grasp of the kind of causal information-preserving property that I have suggested determines the reference of “Paris.” But as this objection comes up so often in discussion, it calls for a more extended discussion.²³

How very young children represent the world to be when they use various words is more a subject for child psychologists than philosophers. All the same, parents and grandparents can have empirical information that bears on the subject, and when they do, it can count *against* the view that very young children refer to the same things we do when they repeat our words. Here is a case in point. One of my young grandchildren has a small dog called “Oskar.” She says “Oskar” when she sees him. I might have thought she was using it as a name for Oskar until I was told that she says “Oskar” in the presence of any small dog. Empirical evidence relating to word usage is very much to the point when considering the question of whether and when children are using a name, or indeed any word, as we do.

²³ My discussion here draws on helpful discussion at the University of Vermont.

We note above the evidence that tells us that we – the mature folk, those of us who know what is going on when we search for the Jonah of the Bible, or Helen of Troy, or Robin Hood, or Bielefeld – are using proper names as parts of information channels. To the extent that some child isn't using a proper name in that kind of way, they are not using the name as we do.

The upshot is that the objection to the description theory from the word usage of very young children faces a destructive dilemma. Either they use, say, "Paris," as we do, in which case we should credit them with a grasp of the information-preserving causal chain, or else they don't use the word as we do, in which case their usage is irrelevant to the right theory of *our* use. Perhaps those who advance the objection from very young children are implicitly supposing that a child's *trying* to use a word as we do is the same as their using a word as we do, but we saw the problems with this kind position earlier (in §4 above).

9. The description theory and interchangeability

I now come to what I heralded as the most influential reason for rejecting the description theory of reference for names. It comes up again and again, in one form or another, in the writings of opponents of the description theory.²⁴

Suppose, runs the objection, that the description theory is correct for proper name "*N*." This means that there is a property, being *D*, that determines the reference of "*N*," and that this is known by competent users of the name. Then "*N*" will refer to *x* if and only if *x* is the *D*, and competent users of "*N*" know this. Then "*N* is *F*" will mean the same as, or be equivalent in some strong sense to, "*The D is N*." But there are decisive arguments against the equivalence of "*N* is *F*" with "*The D is F*," for any "*D*" that might plausibly be the reference-determining property.

The last claim is correct. There *are* decisive arguments against the equivalence. The problem with the objection is that the description theory is not committed to the equivalence. Despite the allure

²⁴ Soames (2002; 2005a) are prime examples.

of the argument to the conclusion that description theorists are committed to the equivalence – indeed, the *a priori* equivalence – of “N is F” with “The D is F,” for some “D,” and despite the fact that description theorists sometimes hold that they are so committed, they aren’t.

The easiest way to see this is to reflect on sentences of the form: “I am F.” There is no mystery about the reference rule for “I” in such sentences. The rule is that it refers to the producer of the sentence. There may, in some cases, be a certain amount of argy-bargy about how to spell out what it takes to be the producer of the sentence. I encourage Fred to run for office. As a result Fred says, when nominations are called, “I am a candidate.” I caused the sentence to be uttered but I am not the producer in the relevant sense. Fred is, and the token of “I” refers to Fred and not to me. Anyone who has wrestled with the problem of deviant causal chains will suspect that it may not be easy to say exactly why, in this example, Fred counts as the producer of the sentence and so as the referent of the token of “I” whereas I don’t.²⁵ All the same, the reference rule is straightforward. “I” refers to the producer of the sentence. What is occasionally tricky is saying why one person counts as the producer whereas another is a causal factor but not the producer.²⁶

This means that a description theory is correct for tokens of “I” in “I am F.” The reference of “I” goes by, is determined by, whatever has the property of being the producer of the token sentence. And that is common knowledge and part of what it is to understand how “I” works. This is fully in accord with our approach to the question, What is a theory of reference a theory of? The information imparted by, how things are being represented to be by, an assertoric use of “I am F,” is that the producer of the sentence is F, and someone who does not know this has a defective grasp of how “I” works in this kind of sentence in English.

²⁵ For evidence for the intractability of saying wherein deviance lies, see, e.g., Davidson (1980).

²⁶ There is also the complication posed by cases where one person produces a token of “I” and another completes the sentence. A reader pointed out that in such cases the “I” may refer to the person who produces the “I,” not the person who completes the sentence. But that is not always the case. Consider a survey asking one to complete “I admire ...” The token of “I” in this case does not refer to the writer of the survey but to the person who completes the sentence.

However, it would be a mistake to hold that “I am *F*” is equivalent to “The producer of ‘I am *F*’ is *F*,” as is widely acknowledged. “I have a beard” is not equivalent to “The producer of ‘I have a beard’ has a beard.” There may be, indeed no doubt are, many producers of the sentence “I have a beard,” and even if there weren’t, someone who says “I have a beard” isn’t saying *inter alia* that there is only one producer of “I have a beard.”²⁷ True, there is only one producer of the token sentence, but that doesn’t help. “The producer of the *token* of ‘I have a beard’ has a beard” faces the same problem. There will be many tokens. It may well be that there is some unique specification of the token in question. But there is no plausibility whatever to the suggestion that “I have a beard” is equivalent to something like “The producer of the token of ‘I have a beard’ at noon, November 15, 2007, in New York, at 365 Fifth Ave, has a beard.” That kind of detail is no part of how I represent things to be when I use “I have a beard.”

It would be even more obviously a mistake to hold that the sentence “Fred believes that I have beard” is equivalent to “Fred believes that the producer of the sentence ‘I have a beard’ has a beard.” Perhaps Fred believes that I have a beard but that I never ever comment on the fact and, moreover, that there is only one person in existence who ever produces the sentence “I have a beard” and that person is, according to Fred, mistaken on the point (they don’t have a beard). In that case “Fred believes that I have beard” is true, whereas “Fred believes that the producer of the sentence ‘I have a beard’ has a beard” is false. Or perhaps I am clean shaven and Fred knows it but thinks (no doubt mistakenly) that there is only one person in existence who ever produces the sentence “I have a beard” and Fred believes this person is bearded. In that case “Fred believes that I have beard” is false, whereas “Fred believes that the producer of the sentence ‘I have a beard’ has a beard” is true.

The moral is that the question of whether the reference of a singular term, “*T*,” goes by a known property, being *D*, when knowledge of this fact is part of understanding the term, is a distinct

²⁷ There are, of course, more reasons than this one to deny the equivalence. When I believe that I have a beard and give voice to it using the sentence “I have a beard,” the content of my belief has a lot in common with the content of my sentence but the belief is not about a sentence. But one decisive reason is enough for our purposes here.

question from the question whether “*T* is *F*” is equivalent to “The *D* is *F*.”

We might have argued that “I have a beard” isn’t equivalent to “The producer of ‘I have a beard’ has a beard” by pointing out that “I” is a rigid designator whereas “The producer of ‘I have a beard’” isn’t, thus reprising an issue that looms large in the debate over the description theory of reference. But that would invite the thought that it might be possible to hold that “I” is equivalent to “the *actual*” This would be a mistaken thought. The reason we gave for holding that “I have a beard” is not equivalent to “The producer of ‘I have a beard’ has a beard” had nothing to do with the difference in respect to rigidity. The message in what is to come in later lectures is that it is a mistake *in principle* to hold that sentences of the form “*N* is *F*” are equivalent to “The *D* is *F*,” and to think one might restore equivalence by rigidifying, or by imposing a special scope reading, is to fail to grasp the fundamental reason why the equivalence fails. All the same, the description theory of reference for proper names is true. The reference of proper names does go by known properties, and knowing what the properties are is part of understanding the role of proper names in representing that the world is thus and so, in passing on or imparting putative information that the world is thus and so.

I am doing a certain amount of ‘territory marking’ here. One might insist that *by definition* the description theory of proper names holds that “*N* is *F*” is equivalent to “The *D* is *F*,” for some *D*, in some strong sense, with special clauses to do with rigidity or scope added. Indeed, critics of the description theory sometimes write as if it was constitutive of the description theory that “*N* is *F*” is *a priori* equivalent to “The *D* is *F*,” for some *D*. If one does so insist, then I am with the majority in repudiating the description theory of reference for proper names. But then I (and we) will need another label for the theory of proper names I defend. I hereby stipulate that what I mean by the description theory is the theory I have been defending above.

10. What is to come and a final objection noted

I can now highlight a major item on the agenda of the lectures to come. I say above that, despite the fact that the reference of proper

names goes by known properties and that knowing what the properties are is part of understanding the role of proper names in representing that the world is thus and so, sentences of the form “*N* is *F*” are not equivalent to sentences of the form “The *D* is *F*,” and that this is so for reasons that have nothing especially to do with the familiar issues about rigidity and scope. How can this be correct? We will be saying how. The story starts in the next lecture, but I will finish this lecture by addressing an objection many will, I suspect, be itching to make around about now.

We noted earlier the distinction between the content a sentence in fact has versus how it comes to have that content, and the correlative distinction between the reference a term in fact has versus how it comes to have that reference. The first concerns how things are being represented to be by a sentence and the contribution a word makes to what we are calling its *ir*-content, whereas the second concerns how the sentence comes to have the *ir*-content it does have and how a word comes to make the contribution that it does make to that *ir*-content. The objection is that when I said that the reference rule for “*I*” in “*I* am *F*” is that “*I*” refers to the producer of the token sentence, what I said was true *only* if I was talking about how “*I*” comes to have the reference it does. In the sense of “reference” tied to the contribution “*I*” makes to how, say, “*I* have a beard” represents things to be, runs the objection, the correct answer concerning the reference of “*I*” is a Millian one. That is, if Frank Jackson says “*I* have a beard,” the content is Frank Jackson’s having a beard – that very person’s having a beard; something we might represent with the ordered couple $\langle Ff, \text{having a beard} \rangle$. In this sense of “content,” the content of “*I* have a beard” varies depending on who says it, and the variation in content is tracked by the reference rule we are discussing.

We will discuss this issue in more detail in Lecture Three, but let me give the short version of my reply here. The informational value of “*I* have a beard” does *not* depend on knowing who said it. Consider an amnesia sufferer waking in hospital with no idea of who they are. On feeling the beard on his face he knows how to say what he has just learnt – use the sentence “*I* have a beard.” It does not matter that he does not know who he is. And someone competent in English who hears the sentence but has no idea who said it (perhaps they are out in the hospital corridor) knows what is been said and

gets information from hearing the sentence. What information? We all know the answer to that question: the information that the person who uttered the sentence has a beard. Ergo, the reference rule we have been talking about delivers reference in the sense of its role in how things are being represented to be, its role of delivering putative information about how things are.