

Guide to Chapter Nine of Gareth Evans' *The Varieties of Reference*

Rick Grush
Department of Philosophy 0119
University of California, San Diego
La Jolla, CA 92093

9.0 What Happens in Chapter Nine

The topic of this chapter is information-based thoughts, and information-invoking singular terms (IISTs). Information-based thoughts were discussed in Chapter 5. They are thoughts employing an Idea whose controlling conception includes information whose source is some object – and if the thought is well-grounded, it will be from the object identified by the thoughts means of object-identification. And information-invoking singular term is a linguistic expression an understanding of which conventionally requires the audience to have an information-based thought, the subject must bring information to bear on interpreting the remark. In Section 9.2 Evans points out that in making a remark, a speaker intends the audience to think of a particular object, and has various resources for helping the audience get to the right object. He distinguishes three things: 1. The referent of an expression as what the public norms would say a use of that expression should refer to in the context of use; 2. The object that the speaker is trying to refer to by using an expression (note that if S word choice is inadequate, then 1 might be different from 2); and 3. The object a speaker refers to in making a remark that uses a referring expression. According to Evans, 1 and 2 must converge on the same object in order for it to be the object described by 3. In 9.3 Evans argues against the position that there is a conventional “referential” use of definite descriptions. The argument is basically that while people can use DDs as referring expressions, people can use lots of things as referring expressions. But none of the arguments establish that in fact there is a conventional meaning of DDs as referring expressions. In Sections 9.4 and 9.5 Evans argues that IISTs are Russellian. In 9.4. the argument is that in order to understand a sentence using an IIST, the audience must not just have information, but must use it intelligently, and this can only be done if the audience believes the object exists. In Section 9.5 the argument is that in order for communication to take place by using an IIST, the object must exist because the object is the only unifying principle behind the way the speaker and audience might (differently) identify the object. In Section 9.6 Evans points out that cases where empty IISTs can be used for apparently fruitful communication when the speaker and audience know the IIST is empty are best understood as involving pretense, which is the topic of Chapter 10.

9.1 Introductory

There are three connected topics in this section. The first is the claim that most uses of referring expressions require, in order for sentences employing them to be understood, that the audience have an information-based thought. There are many varieties of such expressions, including demonstratives and proper names. This is a preliminary to Evans giving a fuller account of this later in the chapter. The second is that Evans contrasts this proposal with a similar proposal by Strawson which maintained that referring expressions required 'identifying knowledge' on the part of the hearer. The problem with Strawson's proposal, though, is that it fails to distinguish identifying content based on information from an object, and identifying content derived from a description. Finally, Evans argues that the role that the information plays is not exhausted by its provision of content, since all the content could be there and the listener might still not know what object the speaker is talking about until the speaker manages to find the right file.

Two things happen in this section. The first is that Evans claims *almost* all referring expressions require, in order to be understood, an information-based thought, in the sense he discussed in Chapter 5. Descriptive names would be an exception. There are several types of expressions he brings up as examples: demonstratives (*this man is F*); past-tense demonstratives (*that man we met last night is F*); testimony demonstratives; and proper names. The doctrine is that in order to understand any of these, the hearer must have a thought employing an Idea whose controlling conception has information from the object. As he puts it:

... the speaker might advert to information he presumes the hearer has from the testimony of others, perhaps from a newspaper article, or a rumour, or a conversation, saying something like 'That mountaineer is F'; here I do not think that the hearer can be said to have understood what the speaker is saying unless he possesses this information and thinks, in a way which is informed by it, 'That mountaineer is F: that's what the speaker is saying.' (VR, p. 306)

Demonstratives are the most straight-forward case because they prototypically don't involve the hearer invoking any prior information from a *previous* encounter or previous testimony. There is information in the CC file, but it is exhausted by the information collected in the current encounter. Of course, if the demonstratively identified object is re-identified as one encountered before, then the files will be merged. The other cases (past-tense demonstratives, and so forth) require the subject to access a pre-existing file as the file for the object the speaker is talking about. The differences concern the nature of the sources of information in that file. It could be general knowledge (which is often the result of testimony of various sorts), or a prior episode in which one encountered the object with the speaker.

The second topic of this section is a discussion of Strawson's views on referring expressions. Evans' point about some expressions requiring an information-based thought has an analogue in Strawson, but Strawson couched the point not in terms of information, but in terms of *identifying knowledge*. And moreover, he was happy to let this apply to objects referred to by description. Understood in this way, "identifying knowledge" is a broader category that includes information, but also includes non-information-based beliefs. For example, beliefs that can be derived from a purely descriptive identification. This might include your belief that Julius is an Englishman, which is not information-based, but does reflect content in your CC file of Julius, whether you had thought about Julius before or not. But Evans thinks that exactly because Strawson's account puts both descriptive thoughts and information-based

thoughts under the general heading of 'identifying knowledge' which Strawson treats as effectively homogenous, it won't do. In particular, it won't do because it invites a conflation of two factors that Evans thinks are related but independent: the mode of identification of the object, and the object as the source of information. Recall, it is only when these two factors converge on the same object that the thought is well-grounded.

The third topic is connected to the first two. It is that the crucial role of the information is not just to provision of information that could be used to construct a description. In some cases the content derived from the information might be part of what allows the subject to achieve discriminating knowledge, but sometimes this is not the case — the discriminating knowledge has another source. For example, with demonstratives, the crucial factor is that the information link allow the subject to locate the object in space. The content delivered by the IL plays little role, and can even be largely erroneous.

Evans' example is a speaker who is attempting to refer to a bird that he saw in the past with the listener. In attempting to jog the hearer's memory, he may provide as much or more content than what is contained in the information the hearer has. But until the right file is retrieved, the speaker can't be said to have understood the sentence. The hearer's latching on to the right file might be accompanied by something like "Oh, *that* bird!" Before the realization of which bird was in question, the listener was in possession of all the content associated with the object. And the crucial moment isn't adding anything to the content, but just the invocation of the right information-based file.

This is an important point that needs to be emphasized. Suppose that the information that the hearer is unable to access at first contains descriptive content $\varphi_1\text{-}\varphi_n$. The speaker might, in trying to jog the listener's memory, mention all of this. Enough so that the listener might genuinely believe, truly, that if what the speaker is saying is true, then there is something that is $\varphi_1\text{-}\varphi_n$ and is F. But notice that the speaker's thought is not an existential thought to the effect that there is something that is $\varphi_1\text{-}\varphi_n$ and is F. And adding a φ_{n+1} to the effect that they both encountered it doesn't help. It's just still a more detailed existential thought. There is still room for a realization on the hearer's part. When the right file is accessed, no new descriptive content is available. The hearer, in accessing the file might, as Evans puts it, say "Oh, you mean *that* bird!" And when that happens, the hearer understands the remark for the first time. And what was needed for understanding was the invocation of the right information from the right source. Call this the *information-invoking thoughts are not existential thoughts* principle. I'll return to this later in this chapter.

Why is it that existential thoughts strongly suggest themselves as replacements for demonstrative thoughts? So much so that the replacement can (according to Evans) be effected without the subject thinking a significant change has occurred? There is a fine line between

- A. Knowing the conditions such that, had they obtained, one would have been in a position to grasp a thought expressed by remark R.
- B. Knowing the conditions that, had they obtained, remark R would be true.

Evans is trying to draw a bold line between these. And a good deal of the resistance to Evans' position comes from not recognizing (for better or worse) this bold line. Evans' position is that for IISTs these are distinct, but almost inescapably confused. And the culprit is existential statements. Surely with a use of an empty IIST of the form *that G is F*, if the existential statement *there is a G right there that is F* there were true, then the speaker would have been in a position to have a demonstrative thought and the IIST would (probably) have worked fine. (I say "might have" because it is of course a possibility that the G that would have made the existential statement true would not be the one intended by the speaker, perhaps because it is possible that a G could be in the vicinity but unseen.) But even though this is a condition such that, had it obtained, a demonstrative thought would have been available, it does not follow that this can be analyzed as the truth condition for the demonstrative thought that would have been enabled.

9.2 How communication is effected

There are two major threads in this section. The first is, as the title suggests, how communication is effected. In particular, how is it that the speaker can get the audience to think of the right object (the speaker's intended referent). There are many resources at the speaker's disposal, including picking a referring expression that places constraints on what the referent could be, exploiting the pre-existing salience of some objects, or exploiting expressions that are conventionally associated with ways in which the audience should think of the referent (which, if the audience employs those ways, should result in the audience thinking of the correct object). The second part is a clarification about how to understand the object that the speaker is intending to refer to. This second bit is tough to state succinctly, so I'll just leave it at that for purposes of this quick gloss.

From the beginning of the section to the bottom third of p. 310, Evans makes a preliminary point that it is not always clear, just on the basis of the grammatical category of a noun (or whatever components of a noun phrase), whether it is information-invoking or not, since many grammatical categories can go either way. Evans gives the example of pronouns, which are sometimes information-invoking and on other occasions not. But most of the time it is understood from context or other clues whether a given use of an *expression* requires that the hearer have an information-based thought or not. Notice that the topic here is natural language expressions and whether they have conventional uses as information-invoking expressions. This is separate from the question whether, on some given occasion, a speaker S might use an expression in a way that it is clear that S intends the audience to interpret the remark in an information-invoking way. Language users often use expressions in a way that coerces them into doing things that go beyond their conventional uses. For example, speakers can often use names as verbs: *That "responsible gun owner" just Second Amended himself in the leg.*

We then get the main topic of this section, which is "how communication is effected". Communication is achieved when the audience gets the right (aka *intended*) interpretation. Exactly what this means is going to be discussed at length in this section, but we get a constraint right away: communication is a mode of transmission of knowledge:

... it is a fundamental, though insufficiently recognized, point that communication is essentially a mode of the transmission of knowledge. In application to the case we are particularly interested in, this means that, if the speaker S has knowledge of x to the effect that it is F, and in consequence utters a sentence in which he refers to x, and says of it that it is F, and if his audience A hears and understands the utterance, and accepts it as true (and there are no defeating conditions), then A himself thereby comes to know of x that it is F. (pp. 310-311)

This knowledge-transmission principle is going to be used in the following way. It might of course be true that on particular occasions language use fails to transmit knowledge. But if a theory of the functioning of language is such that, if that theory were correct, linguistic communication could not reliably serve the purpose of knowledge transmission, then that theory must be wrong.

The first application of that principle is that in order for language to function as a mode of knowledge transmission, the speaker must be able to make it manifest to the audience which object she is speaking of. This is a bigger deal than Evans makes it out, and the discussion here could be clearer. To see the import of this, consider how Evans describes this first application of the principle:

... in order to say something, one must enable an audience to know what it is that one is saying. A speaker who is to say something by uttering a sentence containing a referring expression must make it manifest which object it is that he intends to be speaking about – which object an audience must think of in understanding his remark. (VR p. 311)

Now compare this to a remark Evans made in Chapter 3, after pointing out that expressions of public language have semantic properties independent of what the speaker might intend:

There immediately opens up the possibility of a gap between what a speaker means to say by uttering certain words – what thought he wishes to express – on the one hand, and what he strictly and literally says, according to the conventional meanings of the words he utters, on the other. (VR p. 67)

And he follows up by pointing out that speaker's intentions are irrelevant to determining what was said by an utterance (except insofar as they can tell us which word was intended to be used):

So when we wish to establish what a person is saying in uttering certain words, we must get clear exactly which linguistic counter, so to speak, the speaker is putting forward. Here the speaker's intentions are indeed paramount. But, once it is clear which linguistic counter he is putting forward, the content of what he says is determined by the significance which that counter has in the game, and not by whatever half-baked and ill-informed conception he may have of its meaning. (VR pp. 68-69)

The Chapter 3 point is that the speaker's intentions — the object she is intending to refer to, or the thought she is intending to communicate — do not determine the content of *what is said* (aka “linguistic meaning”, the meaning of expressions in public language as determined by their norms of usage). The Chapter 8 point is that while it is true that the speaker's intentions are not relevant to the *linguistic meaning* of the utterance that is produced, those intentions *are* relevant to determining when communication has been successful. Putting these points together is kind of what is going on in this section. The speaker must, in order to communicate in such a way as to be able to convey knowledge, select expressions such that their publicly established meanings typically have, as their referent, the the

object the speaker is intending to refer to. And since, for Evans, the meaning of an expression is explained as the thought that an audience would have to have to understand it, this means that it must be possible for the speaker to reliably choose expressions that will be such that, in order for the audience to understand them correctly, the audience will think a thought about the very object that the speaker is intending to refer to.

There are many things the speaker can do to try to ensure that the audience identifies the intended object. One is the speaker's choice of referring expression. As Evans puts it, different REs have different "referential features" that place constraints on what the referent could be. This is pretty straightforward. But the key point is that this mechanism is one that exploits a conventional association between expressions in language and required/allowable features of the referent.

Another option open to the speaker is to exploit the salience of objects. Evans doesn't put it this way, but "salience" is in effect a word that encodes the fact that people, including the audience, are probably having their attention, and hence thoughts, drawn to that object already independently of any language use. The speaker can exploit that by using an expression whose conventional meaning is tied to one's attentional state in this way.¹

At the top of p. 313 Evans makes a point that is tied to the knowledge-transmission principle. In order to understand the some utterances, it is not enough that the audience just have any old information-based thought about the correct object (intended referent). Rather, the audience must have the right kind of information-based thought. This requirement applies only to "certain referential remarks" and not all. Proper names have no requirement of the sort, for example. So two questions are: what does he mean by the "right kind" of information based thought?; and why is this a requirement?

Evans' example is a speaker who says "he's had enough" when someone in a line of people faints. The salience of the person falling down would normally be enough to make it clear what the intended referent is. But it is not enough that the audience have any old information-based thought about the referent. The audience must have the right kind. In this case a *demonstrative thought* underwritten by perceptual contact with the referent. Suppose the audience is blind, and does not see anyone fall. But someone leans over and whispers in the listener's ear "Prince Charles just fainted," the subject will have an information-based thought about the intended referent, but won't have understood the original remark. The audience managed to grasp a thought that was the same, or similar, to the one the speaker was trying to convey, but not because the audience understood the speaker's remark. Rather, because she understood someone else's remark.

The knowledge-transmission constraint comes into play. The whispering informant mechanism that happened to give the audience a thought about the intended referent can't be relied upon to function for communicative purposes, and hence is not adequate to count as a mechanism that can reliably serve as a means of knowledge transmission.

Near the top of p. 314 Evans gives another example, "you" and "I". In such cases it is not enough that the audience think of the intended object. They must think of it in the right way. For "you" the

¹ Demonstratives, and their relations to attention are really complicated. But also really fascinating. For a stab at a unified account, see Section 4 of Brovold and Grush (2012).

audience must think of the referent using an “I”-thought. To understand a sentence using “I”, the audience must typically have a demonstrative thought about the speaker. If someone says to me “You are about to step in a puddle”, and the thought I entertain on the basis of that sentence is “the tallest member of the UCSD philosophy faculty is about to step in a puddle” then I have not understood your remark, despite thinking of the intended referent. And this would be true even if I happened to invoke information about the tallest member of the UCSD philosophy faculty – either from testimony, or even from perception, seeing myself in a mirror but not realizing it is me.

So the doctrine is: to understand the sentence, the audience must have the right kind of information-based thought (for certain kinds of referential remarks).

On the passage bridging pages 315-316, Evans puts these points in terms of the Fregean model of communication that he articulated in section 1.5. Here is a take on what Evans is saying that goes a bit beyond what he actually says. The speaker is trying to convey a thought concerning object *a*, the intended referent. The speaker has a variety of resources available in order to get the audience to grasp the thought with the intended referent. Choice of expressions, exploitation of salience, and so forth. Normal proper names are great, because they are just attached to the object, and don't require you to think of the object in any particular way. Another way to get the audience to think of the right object is to influence their thinking, so to speak. I use “that” and get you to use your demonstrative way of thinking. And I do that because I know that we are in a context in which if you employ that way of thinking, you'll think of the referent I want you to think of. Maybe because there is something that is sufficiently salient in our environment, or whatever.

Here's another way to put it. I want you to think of *a*. And for one reason or another I don't use a name for *a*. Here is another option. For whatever reason I know that if you were to employ a specific way of thinking of objects, *w*, then your thought would latch on to *a*. So I use an expression conventionally associated with *w*. You hear this expression, start thinking in manner *w*, and as a result (if all goes to plan) you think of *a*. Here is an analogy. You are blindfolded, and in front of you are a bag of popcorn, and a bag of M&Ms. I want you to grab the M&Ms. You can't see them, so I can't just say “grab the M&Ms.”. But I do know that you have two ways of grabbing bags: reaching out with your right hand, *w_r*; and reaching out with your left hand, *w_l*. And I also can see, from where the M&Ms are located, that if you grab in manner *w_r* you will grasp the M&Ms. So I use an expression conventionally associated with way of grabbing *w_r*. You then reach out in manner *w_r*, and grab the M&Ms. (Yay.)

Note the difference between exploiting “referential features” conventionally associated with expressions, and exploiting “ways of thinking of referents” conventionally associated with expressions (as just described). The referential features (‘he’ must refer to something that is male; ‘you’ refers to the person being addressed, etc.) focus on features of the referent in order to get the audience to get to the intended object (“grab the M&Ms” focuses on the target, not so much on how you get there). But the ways of thinking tactic is largely silent on features of the referent, and focuses rather on the audience's ways of thinking (“grab in manner *w_r*” focuses on what you do, not so much on the target itself).

The claim is that this is a limited vindication of a Fregean model, because in many cases communication does not require that the speaker and audience think of the referent in the same way. For example, in order to understand my sentence “I am hungry”, I am thinking of myself in the particular way that constitutes an I-thought, but if you employed an I-Idea to get at the referent, you’d be thinking of the wrong object — *you*, not *me*. You would need to use a demonstrative Idea targeting me. So the Fregean model, while helpful in some cases and in some respects, doesn’t hold generally.

I should say, though, that it seems that one could articulate a refined version of the Fregean model according to which what is crucial is not that the speaker and audience employ the same way of thinking of the object, but rather that they use ways of thinking of the referent that are conventionally associated with certain expressions in such a way that this association can be exploited to support communication. For example, the speaker can use it to ensure that the audience get to the intended referent. Situations that require the speaker and audience to think *in the same way* would just be a special case. This proposal would not be classically Fregean, in that it would not require that the speaker and audience think of the object in the same way in order to share a thought. But it would be Fregean in that it would have the following feature, which is arguably what was most important for Frege. It would still give a central place, in an account of communication, to language user’s ability to exploit conventional (aka “*objective*”) ways of thinking of objects. At least for many kinds of referring expression. In any case, here is Evans’ summary:

... we can recognize at least a limited applicability of the Fregean notion of sense to expressions of the kind we are concerned with; but without in any way gainsaying the idea that what is primary, for purposes of communication, is the referent. The limited recognition of sense comes in with our claim that understanding the remarks we are concerned with requires not just that the hearer think of the referent, but that he think of it in the right way. But we recognize the primacy of the referent by recognizing that the hearer always confronts just one question, ‘Which object does the speaker mean?’ — not two questions, ‘Which object does the speaker mean?’ and ‘How am I intended to think of it?’ The second question is answered in passing; for if he understands the remark, he will know which object is meant; and in the normal course of events (i.e. without assistance from others, etc.), he will know which object is meant only if he thinks of it in the particular way intended by the speaker. (pp. 315-316)

That was the first major thread of the section. The second major thread is getting clear on what counts as the intended referent. The obvious suggestion is this: the speaker is entertaining a thought and that is the thought the speaker wishes to convey. In grasping the thought, the thinker is thinking of some object. As Evans puts it:

...in the normal case, the speaker will make information-invoking uses of referring expressions in order to give expression to his own information-based thoughts, and his intention could itself be expressed in one such thought: ‘By t I mean that one.’ Now, it would be natural to take the intended referent of this kind of use of a singular term to be the object which these thoughts concern; but this natural position cannot be accepted as it stands. (p. 316)

Why not? The problem is that the speaker’s conception of the referent may be quite complex, and may rest on a variety of identifications (or mis-identifications) of the object. Evans example is that I met a person, *b*, yesterday and take myself to be encountering them again today -- though in fact it is a

different person, *a*. If one tries to excogitate from my thoughts which object I intend to refer to, we might be in a muddle.

While my thoughts might be complex in this way, language is generally geared towards the task of favoring non-compound ways of identifying a referent. As Evans puts it:

There seem to be some remarks that the subject might make in this situation of which *a* is clearly the intended referent, and other remarks of which *b* is equally clearly the intended referent. For example, if the speaker uttered the words 'That man over there is F' the intended referent would surely be *a*, whereas if he uttered the sentence 'That man we met yesterday is F' the intended referent would surely be *b*. (VR p. 317)

Evans makes appeal to what he calls the lowest level action plan. Even if my Idea of the object is complex, and perhaps erroneous in some ways, it is still true that my remark will latch on to one way of identifying the object, and this is the intended referent. Evans' provides the following analogy:

Thus the notion of the intended referent is rather like the notion of a target. Suppose the subject, in the case we have been considering, had aimed a gun at the man he could see. Even if his general plan was to shoot *b* — for example, because the offence he wished to avenge occurred in the previous encounter — it is undeniable that *a* was his target, and that he intended to shoot *a*. His lowest-level action plan concerned *a*; success in it would involve the shooting of *a*. Similarly, *a* is the speaker's linguistic target when he utters the sentence 'That man over there is F'; this time he is directing, not a gun, but his audience's attention. (VR p. 317)

So the speaker's intended referent is the object (if any) that figures in the speaker's lowest level action plan.

The next topic concerns not what the intended referent is, but what object (if any) the speaker refers to by *making a remark*. The fact that *a* is my intended referent does not guarantee that, in making a remark, I successfully refer to *a*. Evans says that it is also necessary that I make my intentions manifest, and (more importantly) that I choose a linguistic expression whose conventional interpretation would establish *a* as its referent. For example, even if my intended referent is definitely Arthur C. Clarke, I won't have referred to Arthur C. Clarke with my remark if my remark is "Barack Obama is famous", for despite my intentions, I will have used an expression that is not "conventionally suitable for reference to the intended object" (p. 319). So notice that three things are being distinguished:

1. The conventional meaning of a referring expression as a word in public language.
2. The speaker's intended referent. (This is to be understood in the 'lowest level action plan' sense.)
3. The object a speaker refers to in making a remark that uses a referring expression.

The doctrine is that in order for (3) to be *a*, it is necessary that both (1) and (2) converge on *a*. Evans makes three clarifications. The first is one that has already been accounted for in my remarks above. It is that while (2) is necessary for (3), it is not sufficient.

In a passage bridging the bottom third of p. 319 and the top third of p. 320, Evans argues against the position to the effect that (2) is not necessary for (3); in other words, that (1) is sufficient. The argument is this: uses of referring expressions are often ambiguous in that the expression by itself underdetermines the correct interpretation. For example, when I say “That kid is going nuts” in a room with several energetic children. We are, *ex hypothesi*, ruling out speaker’s intentions in doing the disambiguating. The only other option is supplementing the conventional meaning of an expression with various contextual clues. Evans has two arguments against these contextual clues being sufficient, apart from speaker’s intention, to establish the referent of the expression. The first is that this principle still seems to require that the target is the speaker’s intended referent, and so why not just include it up front? The second argument appeals to the knowledge-transmission principle. If this proposal were correct, then it would be a built in structural part of language that an audience could be *correct* in identifying, as the referent of an expression (used on a particular occasion by a particular speaker), some object other than the one intended by the speaker. And if this were the case, then language would not be suitable for the reliable transmission of knowledge.

Note that Evans is not denying that the audience must use contextual clues. Of course they must. The question is whether, when those clues manage to implicate an object *a* that is different from the object the speaker intended, *b*, is *a* the object the speaker referred to in using that expression ((3) above)? Evans plausibly denies this.

The final paragraph and its associated footnote attempt to clarify some terminology. The last paragraph distinguishes:

... the referent of an expression as used by a particular speaker on a particular occasion; that of the intended referent of an expression; and that of the object which the speaker means. (p. 320)

The first is (3) above. The second is (2). The third (which object the speaker ‘means’), is the object, as the speaker conceives of it, that is the object of the lowest level action plan. This would include, as components of the speaker’s idea-of-the-object re-identifications that might in fact be erroneous, but in any case would be beyond the object, narrowly conceived, that is the object of the speaker’s lowest level action plan. The footnote adds a fourth category: the object that would satisfy the speaker’s high-level goals, in the sense of the psychological idioms discussed in Section 5.3.

9.3 The information-invoking use of definite descriptions

In this section Evans will argue that the considerations that philosophers have given in support of the idea that there is a “referential use” of definite descriptions do not in fact establish that conclusion. The key point is that speaker can “use” DDs in a manner that clearly indicates that the speaker is intending to invoke an information-based thought in the audience. Evans claims that this is what proponents of the “referential use” of DDs base their case on. But this fact is insufficient to establish the conclusion that the grammatical category of DDs is semantically ambiguous, containing two semantic kinds: pure DDs and referential DDs. English allows people to use expressions in a ways that

subvert their conventional uses. Names can be use as verbs; and even existential expressions can be used in an information-invoking way. Evans argues that when one looks at the considerations that would actually establish the semantic ambiguity, it looks like DDs are not actually ambiguous.

The topic of this section is what is commonly known as the *referential use* of definite descriptions. The standard example is: you are at a party, and upon seeing someone drinking from a champagne glass, you tell your interlocutor “The man drinking champagne is a spy,” or something of the sort. The claim is that in this case the NP “the man drinking champagne” refers to the person you are looking at and who prompted the remark, even if that person is a woman, and is drinking water from a champagne glass — that is, even if they do not fit the description. And the is true even if there in fact is, out of sight perhaps, a man drinking champagne who is a spy. The NP would refer not to the person who fit the description, but to someone who it doesn't fit.

Evans prefers the term “information-invoking use of definite descriptions” because he thinks that it is the invocation of pre-existing information about the referent that is key, not the fact that the description is referential. He has tried to argue that descriptive names are genuine referring expressions, but clearly that is not the sort of thing that the proponents of the ‘referential use’ of definite descriptions has in mind.

Evans points out that everyone agrees that there are *uses* of DDs that fit the example. But there are two possible explanations:

1. The grammatical category of DDs is ambiguous between two genuine (fully conventional) semantic kinds: pure DDs (which Evans and almost everyone agrees are best treated as quantifiers) and something like “referential” DDs.

2. The grammatical category of DDs is unambiguous, they are all quantifiers. But in some cases people can employ them in such a way as to clearly intend that they be information-invoking.

Evans motivates the claim that the second option is a live option by pointing out that one might utter the sentence “Under that tree stands an honest man” and clearly intend it to be a referential use directed at someone under a tree in the shared perceptual environment. But this by itself wouldn't be enough to license the claim that “a man” is semantically ambiguous. It would just mean that, like many other expressions, context can override their usual standard conventionalized meanings.

So now the question is how might one argue that they are genuinely semantically ambiguous? Evans discusses the general form that such an argument could take, which would be to show that

A) There are situations in which the truth conditions for the two readings (quantificational and referential) diverge,

and

B) Our intuitive judgments “of correctness or incorrectness” support the referential interpretation.

Here is what (B) is getting at. If the goal is to argue that there is a fully conventional information-invoking use of the term, then that is a claim that goes beyond mere truth conditions, and concerns how expressions are *correctly* used. I might say “me hungry now”, and it would be clear what I am saying, and might even be *true*. But despite its truth, in uttering this sentence I will have violated grammatical principles. My statement while true, would not have correctly followed the conventions of language.

Top of 322 To the bottom of p. 323, He then points out that the typical sort of argument that gets provided for there being a genuine ambiguity envisions a situation in which someone says “The man drinking champagne is F”, manifesting an intention to talk about a woman who is not drinking champagne. Evans' point is that while in some sense it is obvious that we would say that the person is the referent of the expression, it is also clear that we would not judge the remark to be correct:

We would not judge an utterance of ‘The man drinking champagne is F’ to be correct in the circumstances envisaged above. Undeniably, a mistake has been made, and the sentence should not have been uttered. (VR p.323)

The point here is that if there is a fully conventional use of DDs as referring expressions, and their truth conditions were such that the φ is F can be true even when there is nothing that is both φ and F, then there shouldn't be any intuitions to the effect that the remark was incorrect if that expression was used in those conditions. By way of comparison, the word ‘bank’ is multiply ambiguous in English. It can mean the side of a river, a financial institution, among other things. Because the term is genuinely ambiguous with fully conventional uses, none of these invoke any sense of a mistake of any sort having been made:

1. The bank was closed on the day I went to deposit the rent money.
2. The bank is steep and slippery.

But since “the man drinking champagne is F” would not be correct (a mistake has been made”) in the situation envisaged, Evans claims that this argument does not establish that the referential use is a legitimate conventional sense of the term.

From the bottom of p. 323 to the middle of p. 324 Evans explores a different possibility. But before he gets into that example, he establishes a background claim. Evans points out that there is a sort of demonstrative expression, what he calls a ‘testimony demonstrative’, that has a clearly conventionalized information invoking sense. His example is a speaker making reference to a widely shared story about the pope being assassinated, and saying “That assassin of the Pope is Italian.” The ‘that NP’ is what Evans calls a *testimony demonstrative*. Its conventional meaning requires that it be interpreted in terms of information (from testimony, news stories, etc.). And Evans points out that it wouldn't be correct to take “that NP” as a pure description. Evans underscores this via the following scenario: the widely disseminated story about the Pope is a fabrication. However, unbeknownst to anyone but a handful of the Pope's closest associates, the Pope actually has been assassinated. And let us assume that this real assassin is Italian. In such a situation, Evans claims that someone who said:

That assassin of the Pope is Italian.

averting to the shared (mis)-information about the Pope, would not have said something either true or correct. This is because the 'that NP' is specifying that the NP must be the one that is the source of the information that is being appealed to by the remark. So the point is: with an expression that is known to be information-invoking (referential), 'That φ is F' is not correct if the putative source does not exist, even if it turns out that there is something that is φ and F.

OK, with that background in mind, Evans explores a speaker who, in a similar situation, uses "the φ is F", that is, a referential (information-invoking use) of *the φ* . Evans claims that in the situation envisaged (widely disseminated mis-information about the Pope being assassinated by an Italian, while the Pope has actually secretly been assassinated by an Italian), a speaker who said:

The assassin of the Pope is Italian.

would have said something true, and the remark would not plausibly be regarded as incorrect — even if the usual contextual clues were around to indicate that someone intended to be referencing the shared information. This is another knock against the claim that there is a fully conventional information-invoking use of expressions of the form *the φ* .

From the middle of p. 324 to top of p. 325 Evans discusses another potential argument that might be used by the ambiguity proponent. This is that speakers who use the φ referentially clearly don't mean to be implying that φ is uniquely instanced. A defender of the non-ambiguity view might reply that it is still a definite description, but there is an implied supplementation to the description that would make it uniquely identifying. Evans admits that this might work in some cases, but there are other where this response seems inadequate, where it seems most plausible to say that what the speaker is doing is trying to find the right object (the one the speaker is implying is the information source) as opposed to the right descriptive supplementation.

Evans summarizes by saying that the important point isn't whether or not the ambiguity view is correct. Rather, what is important is to understand the thesis. What it would mean for there to be such a genuine ambiguity in the grammatical type *the φ* between a purely descriptive *the φ* , and an information-invoking *the φ* ? And Evans claims that the current (at his time of writing, obviously) literature has not been clear on this, since it has taken evidence that really only shows that speakers can *use* the NP in an information invoking way to argue that the NP is semantically ambiguous between two genuine (fully conventional) semantic types. Recall his example "Beneath that tree stands an honest man."

The upshot is that all that has been established by the current literature is that there are situations in which the speaker uses *the φ* when they clearly have a particular object in mind. But this is not sufficient to show that *the φ* is ambiguous in the advertised way. That argument, if sound, would be sufficient to show that many categories of NP are ambiguous in that way, when they are clearly not. Moreover, the arguments in favor of the ambiguity view don't work, because they fail to show that the interpretation of the expressions as quantifiers is incorrect even in the situations envisaged, which it would have to be if there were a fully conventionalized sense of *the φ* in that way.

9.4 Information-invoking singular terms are Russellian

This section is one extended argument to the effect that information-invoking singular terms (IISTs) are Russellian. The sequence of steps in the argument is: In order to understand an IIST (to know what the speaker is saying), one must use information one has concerning the (purported) referent; in order to use information in the required way, a subject must believe that the object exists; the belief can't be a false belief, since it can't be the case that one must have a false belief in order to grasp a true proposition; so the belief must be true.

IISTs include demonstratives, past-tense demonstratives, testimony demonstratives, and others. And we are considering the situation where the audience does have some information. This could be because both parties can see the demonstratively identified object, or because both are taken in by the same illusion, or because the speaker is subject to an illusion and the audience, while realizing it is an illusion, can make out the features that are supporting the illusion — the curved branch that looks like a mountain lion's tail, for instance. If the audience has no information, then they will presumably interpret the remark as best they can by employing an existential thought of some sort: *that mountain lion looks dangerous* would be “understood” by the audience who can't see a mountain lion along the lines of *there is a mountain lion there that looks dangerous*.

I use scare-quotes because in falling back to an existential thought, the audience is not actually understanding the original remark. Rather, they are grasping a thought that is as close as they can get to the sort of thought that would have been required to understand the remark. Recall the *information-invoking thoughts are not existential thoughts* principle I laid out in my discussion of Section 9.1.

IISTs require, for their understanding, that the audience have some information that they bring to bear on their interpretation of the remark that employs the term. Evans proceeds by first establishing a lemma to the effect that in order for the audience to invoke this information in a way that could constitute understanding the remark, the audience must believe that the object exists. Evans of course agrees that a subject can *have* information in absence of an object, and can also have information despite not believing the object exists. The question is, can a subject *use* that information in an attempt to understand a remark without believing the object exists? (Note that the next chapter is an extended discussion of *pretense*, and in particular how people can *pretend* that there is an object in order to grasp pretend thoughts. A full appreciation of Evans' argument in this chapter requires seeing why, on his account, people have the intuition that they can perfectly well understand such expressions — they have automatically switched into pretense mode without really being fully cognizant of it. See Chapter 10 for details.)

The argument for this lemma runs from the bottom of p. 326 to the top of p. 330. The first premise is that invoking the information — bringing the information to bear on evaluating the remark — just means that the audience is applying the information in assessing the truth, falsity, plausibility, etc., of the remark. The important point is that the information can't just be hanging around in the vicinity or be used willy-nilly. Understanding an IIST requires the audience to intelligently engage their information in their appreciation and evaluation of the remark. A listener who did not do this would obviously be a seriously crippled language user, one who would gain little from conversation employing IISTs. Suppose you have information about your neighbor's dog that it is aggressive and dangerous. I

tell you “that dog next door jumped the fence and is playing with your young nephew.” Presumably you would, using your information intelligently, realize that your nephew is in danger. Imagine someone who did not bring their information intelligently to bear in their interpretation of IISTs and you'll quickly agree with the point Evans is making here.

So the question is, what is the *justification* for bringing this information to bear in this way? There are three options. The first Evans does not discuss explicitly, which is that there is no justification. In such a case, bringing the information to bear with no justification would be what one does in pretending that there is an object. It is to act *as if* there were a justification. (See Chapter 10.)

The second option is that the justification is that the information is shared with the speaker, and the speaker intends this shared information to be brought to bear in the interpretation of the remark. Evans illustrates this first with an analogy to understanding a descriptive name. A speaker says *Julius is F*, and in understanding the remark the audience has the thought that the φ is *F*. But why does the audience bring “information” to the effect that Julius is *the* φ to bear in this case? (The scare-quotes are because it isn't really information in this case, more like *content*). And the justification for bringing this content to bear is that, in using the name “Julius”, the speaker is making it manifest that she intends *this content* to be used in appreciating the remark. And this can be done because both parties know that the name was introduced in terms of that descriptive content.

Now what about an IIST? In this case the content will be genuine information, and we can represent the content of this information as φ_1 - φ_n . The justification for invoking this information in appreciating the remark cannot be that it is the speaker's intent to say of something that it is φ_1 - φ_n and is also *F*, since there is no guarantee that the speaker will have the same information. It would be a coincidence if the speaker does, and a lucky guess if the speaker happened to believe it when it happened to be true. (And a mechanism that relied on coincidence and luck to communicate would not be up to the task of transmitting knowledge.)

The third option is that the audience's justification for bringing the information to bear is that the speaker believes that this information to be *of* the object, and the affiliated content to reflect the features of the object. In this case, the audience's goal is to use the speaker's remark that “a is *F*” to get at the right object *a*, and then appreciate the content of the remark as saying, of an object that is φ_1 - φ_n , that it is *F*.

The only possible justification of the belief that, if what the speaker said is true there is something which is φ_1 - φ_n and *F* is that it follows from some belief of the form ‘The speaker is referring to *a*’, together with a view as to how things stand with *a*. (VR p. 329)

So, even though it is possible to possess information in the absence of belief, it does not appear to be possible to bring this information to bear, coherently, upon the interpretation of a referential remark, unless one believes that there is a particular object to which the speaker is referring, and to which the information concerned is faithful. (One can go through the motions of bringing information to bear if one pretends that there is a relevant particular object: see chapter 10. But we can leave ‘conniving’ uses of singular terms on one side for the moment.) (VR p. 330)

From the middle of p. 330 to the top of p. 331, Evans handles an objection to his view to the effect that all it shows is that the audience must *believe* that the object exists, not that the object *actually* exist — as would be required by RP. For so long as the subject *believes* that the object exists, then the subject will have a justification for invoking the information in a genuine (not pretend) manner.

Evans' argument against this is a bit odd, but maybe it has force. The argument in its entirety is:

... if there really is something said, then it cannot possibly require a false belief for one to know what is said. Truth is seamless; there can be no truth which it requires acceptance of a falsehood to appreciate. (It is surely an axiom of any theory of truth that any true proposition is compatible with any other true proposition.) So the purported 'understanding' is not the knowledge of any truth about what is said; and hence it cannot be genuine understanding at all. (VR, p. 331)

Another premise is being used here, one that has come up before but not so far in this chapter. It is that understanding a remark is knowing what is said by it. This isn't restricted to IIST. The objection is basically that so long as the audience believes there is an object, then there is something said by the remark (the empty IIST), and the audience can know that this is what is said (since this is what understanding the remark would be). The response is that it can't be the case that there are truths that one can't know unless one also believes a falsehood.

This might not carry full conviction with every reader. But the next section provides supplemental arguments.

9.5 Supplementary Arguments

This section produces arguments for the Russellian status of IISTs that differ from that of the previous section. The argument of the previous section basically was that in order for an audience to understand an IIST, they must believe the object exists, and that belief must be true. The arguments here hinge on the idea that communication requires a convergence, of some sort or other, between the speaker and audience, and the only thing generally capable of serving that role is the object. The first argument is based on the fact that communication requires that information be passed, and given that the speaker and audience may have very different information about the object, the only unifying principle would seem to be the object itself. The second argument grants (for sake of argument) that the speaker and audience associate the same information with the expression. But communication can't occur if the information links don't converge. And there are only two ways they might converge: the same object is the source, or both information streams come from the same non-object source. The first option yields the conclusion that IISTs are Russellian. The second is dealt with in Chapter 10.

Evans starts by noting a truism:

To understand a remark made with the use of an information-invoking singular term, it is obviously not sufficient merely that one make some judgement of the general kind we have indicated; understanding, as we would naturally put it, requires that one make a judgement of this

kind about the right object — i.e. about the object which the speaker refers to with his use of the term. (VR. P. 332)

The point could be made even more generally (applying to any singular term, not just IISTs): in order to understand a speaker's remark about an object, the audience must make a judgment about *that* object. If the audience made a judgment about a different object, or no object, then it would be difficult to see how they could have understood the speaker's remark. Evans realizes that putting things this way guarantees that IISTs will be Russellian. And in this section he argues that there is no other way to "state the requirements for understanding" such an utterance.

The first suggestion concerning what might underwrite communication if not the object is the way in which the speaker and audience (purport to) identify the object. If they have this in common, then the suggestion is that this can underwrite communication even if there is no object. Evans uses demonstrative IISTs as an example. In Part II of the book Evans has argued that demonstrative thoughts are Russellian, and so if the conclusions of Part II are accepted then his argument will be short. But he will proceed here in a manner more friendly to his opponent, by ignoring Part II, and assuming that demonstrative identification can be explained as a type of description (since a description seems to have the highest chance of being the kind of thing a speaker and audience could share in absence of an object).

The suggestion then would be that a demonstrative thought would be of the form *the G that is at P*, that is, something like *the object/fish/ship that is right there*. This content, it would seem, could be graspable by both speaker and audience even if there is no G at P. But Evans points out that demonstratives can support communication even when S and A are discussing an object big enough such that their contact with it is at different locations, or through different modalities. The former would be a situation in which A's description would be *the G that is (partially) at P₁*, and A's description would be *the G that is (partially) at P₂*; the latter case would be a situation in which A's description would be the person whose face looks like φ_1 and A's description would be something like the person whose voice sounds like φ_1 . Because nothing in the descriptions requires that the G (or the person) identified by A's description be the same one identified by A's description, this approach cannot explain how these expressions can underwrite communication:

Even when descriptive identification is admittedly at issue, it cannot be generally required for communication that speaker and hearer identify the object by reference to exactly the same descriptions. (VR. p. 334)

From the middle of p. 334 to the bottom of p. 335 Evans addresses the fact that his argument here might seem to show, at best, that identical descriptions are not *necessary* for communication, not that they aren't *sufficient*. And in order to serve his purpose, he would have to show that they aren't sufficient. In footnote 34 (the brackets indicate that McDowell is more author rather than passive editor here), points out that *given* that S and A employing identical descriptions is not necessary for communication, and so:

...the only thing that can unify the quite diverse informational relations this allows us to recognize as possibly obtaining between the information of a speaker and a hearer who communicate in the way we are considering is the existence of an object; so we see that that is a necessary condition for the occurrence of such communication. (VR, p.334, n34, emphasis added)

The argument here is to move from what might be thought to be sufficient *in some cases*, to what sort of consideration would be necessary given the range of descriptions that might possibly be employed by S and A. The move here is one of uniformity of explanation: we need a mechanism that will work in all cases. While the special case of S and A using the same description might be theoretically adequate in some situations, it isn't a mechanism that can explain why communication happens in the wide range of cases in which it actually happens, since in many of these cases the descriptions are not the same. And so we are led again to recognizing the object itself as the focal point:

This argument, then, which we might call the argument from diversity, attempts to show that the only candidate communication-allowing relation, between the thoughts of speaker and hearer, which is discoverable in the absence of an object — that they both involve exactly the same way of identifying, or purporting to identify, an object — is far too strong a requirement to impose upon referential communication in general. (VR, p. 335-6, emphasis added)

So that was the first argument. The second one starts by allowing Evans' opponent the premise that S and A use the same description to identify the referent. Evans points out though that this is not enough, since their sources for this common content involved in *the φ* might be different. S may have heard about *the φ* from testimony tracing back to a certain person's doings, while A's information about *the φ* derived from someone fabricating a story. In such a case, though the descriptions would be the same, A would not be thinking about the right object (or any object — A's thought would be ill-grounded in the sense described in Chapter 5). And so communication would not take place. In order to isolate those situations in which communication can take place, reference would need to be made to the "origin or pedigree of the information upon which speaker and hearer respectively base their thoughts" (p. 336). That is, the object.

And it can't be the source (in what Evans calls the "journalistic" sense of the word, that they heard the testimony from the same person), because it is possible that people can communicate even when they have different sources (in this sense of source), so long as the information ultimately traces back to the same object:

... a speaker and a hearer may attend a party at which they circulate independently of one another. On the next day, the speaker may say 'Did you meet that Russian? He was extraordinarily drunk', and the hearer may understand him in virtue of retaining and invoking information acquired in quite separate episodes from those which gave rise to the speaker's information. This constitutes communication because, and only because, the same object was involved in both sets of episodes (VR., p. 337)

9.6 Understanding without an object?

This section is setting up Chapter 10. Evans divides the situations in which empty IISTs are used into two kinds of case. The first sort of cases are where A has no information from the putative object. These are discussed in p. 338. For example S is hallucinating, and says "He's coming to get me!" Of this situation Evans says:

...I expect everyone to agree that he has failed to say anything (for he has clearly failed to make manifest the sort of intentions he would need to have made manifest in order to have said something). Nor do I expect total opposition to the view that there was not anything that he meant to say. I expect people to agree, in other words, that this case is quite different from the case of someone who points into empty space because he is wearing inverting prisms. For it is extremely difficult to see what belief available to the hearer could count as knowledge of what the hallucinating subject even meant to say. (VR., p. 338, emphasis added)

It is important to remember again that Evans will admit that there are available to A thoughts very similar to those that would have been available to A had S not been hallucinating. In the veridical situation A would have to have a demonstrative thought about *that man*, and while this is not available in the non-veridical case, A can plausibly suppose that some sort of existential thought along the lines of *there is a man there who is coming to get me!* That is close, in some respects, to the thought A would have had had there been an object there. But the point is that while A can grasp this, it must be recognized that this is not the thought S intended to communicate. S wasn't trying to make an existential statement. Recall the *information-based thoughts are not existential thoughts* principle I discussed in my remarks on Section 9.1 above.

From the top of p. 339 to the end of the section Evans discusses the sort of case that he thinks will motivate people to question his program, where there is a "background of shared information":

Perhaps speaker and hearer take the same stone object in the shadows to be a woman, or perhaps they have both heard a story from the same source. When there is this common information, or common misinformation, it seems very tempting to suppose that we can draw just the same distinction between understanding and misunderstanding as we do in the case where there is a referent. We cannot say that the hearer who understands has got hold of the right object; but he can, so to speak, attach the remark to the right information. (VR p. 339)

He points out that he's argued in this chapter that when S or A don't realize that the IIST is empty then communication can't happen. But what about situations when they know that their information is of no object? It seems that in such situations interlocutors can engage in meaningful conversation. Evans admits that something very much like meaningful discourse can happen in this circumstance, and it falls under the heading of what he will call, in the next chapter pretense.

9.A1 Appendix Section 1

This section of the appendix is straight-forward. One way a speaker might try to get an audience to think of the intended object is by putting descriptive material in the IIST that S presumes A will associate with the intended object because S believes that that information is true of the object. S says "That tall Russian at the party last night was hilarious." S believes that the person is tall and Russian, and partially because of this believes that A, who was also at the party, will believe of the same person that they are tall and Russian. (Here this is called "truth-based" because the source of S's confidence that the information will be shared with A is that the descriptive content truly describes the person.)

A different sort of case is one in which S doesn't believe that the descriptive content truly describes the person, but nevertheless is confident that the descriptive content will allow the speaker to get to the right object. There are actually (at least) two cases of this, though Evans only talks about one. The two cases are S knows that A mistakenly believes that a is *the* φ ; or S knows that while A does not believe that a is *the* φ , S knows that A knows that φ is commonly enough associated with a that A will still get to the right object. This would usually constitute something of a joke between S and A. For instance, if S and A are politically liberal, but are at a gathering of politically conservative friends, they might talk about Obama using an expression like "Our Constitution flouting President."

9.A2 Appendix Section 2

This short section is little more than a gesture in a potentially interesting direction. The background is the Fregean model of communication outlined in Section 1.9, which hinged on speaker and audience having the same thought in order to communicate. Evans has already pointed out that while this seems to work fine for the case of demonstratives, it seems to be completely unworkable in other cases, such as 'I' and 'you'. The current topic is thoughts of the sort that would be potentially amenable to the treatment Evans gave of 'today' in his discussion of "dynamic Fregean thoughts" in Section 1 of the Appendix to Chapter 6.

Recall that the idea was introduced with the following suggestion: yesterday I had a thought that I expressed with the words "Today is fine"; and today, I might grasp that *very same thought*, though I would express it with "yesterday was fine." That is, on both days I am thinking the same thought, about the same day, and this is explained by a capacity I have for tracking my location in time. Similar remarks would hold for a thought about location x I would express as "here it is warm." When I move to a different location, I could express the same thought with "there it is warm".

The suggestion is that this capacity might be linked to an ability of S and A to understand each other when S says "It is warm here" and A understands that with a thought that would be expressed as "It is warm there."

But this suggestion isn't developed here at all.