

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

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

5.0 What Happens in Chapter Five

Chapter 4 discussed Ideas of objects, the topic being to get clear on how thinkers distinguish the objects of thought from all others (so they satisfy RP). In Chapter 5, the topic is information. We collect, retain, and distribute information about objects, and often (attempt to) think about these objects. So stated, in any given case, it is an open question whether, in any given case the object that is distinguished by the Idea employed is the same object that is in fact the source of information. When I see a person in the corner, and hear them talking (saying something like 'I am thirsty'), I identify them by their spatial location, and the person identified in this way is the source of the information I am using when I think that person is thirsty. But if I made a mistake — suppose it was really a nearby ventriloquist who said "I am thirsty" — then my Idea would be identifying a person (the one I standing in the corner) who was not really the source of the information.

Evans claims that for an information-based thought to be well-grounded, the object identified by the Idea must be the object that is also the causal source of the information. This makes casual contact with the object a

necessary condition for (information-based) thought about the object. But it is not sufficient: one must still have discriminating knowledge to satisfy RP.

Evans uses this fact — that it is our goal to think about objects that are the causal source of the information we have — to explain why we find it natural to attribute to people beliefs about objects they have information from even when they don't have discriminating knowledge. For instance, we find it natural to say that the subject is thinking about the steel ball from which her information derives. We often attribute to people intentional states that concern that objects that satisfy their high-level goals even when strictly speaking they have no such intentional states. A key example is that, if you know I have a high-level goal of registering a complaint, you might say “You want to see Mr X”, even if I've never heard of Mr. X.

5.1 Information-based thoughts: Introductory

In this section, Evans introduces terminology for two notions. I'll discuss these in the opposite order. The second is 'controlling conception', which is a mental file, a sort of dossier, of content/information (including misinformation) we have about an object. They can sometimes be sparse, e.g. for objects we have just encountered for the first time; or quite extensive for objects that we have seen or heard a lot about. The first is that of being an 'information-based thought'. These are thoughts about objects whose controlling conception includes information (including inaccurate information) whose causal source is (presumed to be) the object itself.

This section is relatively clear and straightforward. It involves mostly introducing some terminology for some uncontroversial observations. The two main notions are that of a *controlling conception*, and of an *information-based thought*. I will discuss them in the reverse order Evans does.

A *controlling conception* is essentially a mental file, or *dossier*, of content that is collected together because it is taken to concern a single object (and typically it does, but cases in which it is taken to be information from one object but really isn't will be discussed). To take as an example a minimal case: I might begin to speculate about the world's tallest spy, and in such a case I would create a new mental file for this person. Initially the file's contents would be limited to the content of the stipulation: this person is a spy, is the tallest current spy. Other content might get into the file as a result of inference (this person is probably older than 5) or prejudice (this person is probably a man).

I might also start a file for an object that I encounter perceptually, for example, a hawk I see in a nearby tree. In that case, the content collected will be *information*, since it derives from causal contact with the object — as opposed to the spy example above, where the content was not information, because I didn't get it through casual contact. Evans will discuss what he means by information in the next section. The dossier contents will initially be limited to what I get from my perceptual contact, including its various perceivable properties (it is brown-ish and bird-shaped), and particulars of my encounter with it (I first saw it yesterday down by the river).

These controlling conceptions can become quite extensive, even encyclopedic, in scope. The file associated with each of your family members for example will contain information, including misinformation, from thousands of extended encounters, and other content inferred, misremembered, derived from rumor, or whatever.

Files are not individuated by the objects they concern. Rather they are individuated on psychological grounds. Suppose I return to the parking lot after grocery shopping and see a dent in my car, marked with grey paint scrapes. In this case I might create a file for the person who dented my car (they drive a grey vehicle, they were in the vicinity of location L on date D); quite independent of this, I might have a distinct file for my friend's mother (she has red hair, she makes ridiculous cinnamon rolls). And these will be distinct files in my cognitive system even if the person who dented my car is my friend's mother. Of course, if I learn about the identity, I can merge the content of the files.

This leads to the second of the two concepts introduced in this section, which is *information-based thoughts*. Evans' use of 'information' is not the same as that used in the mathematical field of information theory. For Evans, information is content whose causal source is some object, and crucially, it can be incorrect — mis-information counts as information if it derives from the object. A few examples might help. When I begin to think of the world's tallest spy, the specifications in the file to the effect that this person is a spy, that they are older than 5, and so forth, were not derived from any causal link to the object. They were the result of stipulation and inference. On the other hand, if I am currently looking at an object, my file for it might include its color, shape and other content that are derived from my perceptual contact with it. Three sources of information that Evans will discuss are perception, memory, and testimony.

Evans introduces a final bit of terminology, which is what he calls a 'pure case'. This will be a situation in which a thought employs a controlling conception (file) that is newly minted, and contains only the content that triggered its creation. This could be a file opened when I first encounter an object perceptually, or when a friend I am conversing with begins talking

about someone they saw on a bus, or when I start to speculate about 'the tallest spy'. It rules out situations in which a file contains information from multiple information-gathering episodes.

5.2 The informational system

In this section Evans has two main topics. The first concerns the nature of the information that is gathered through perception and testimony. The main point is that a specification of the informational content need not (though in some cases it might) make reference to the particular object that was the causal source of the information. He gives as an example an analysis of the content of a photograph, such that two photographs of different objects, but visually indistinguishable, can have the same content.

The second is what he calls the information system, which is a set of mechanisms – both intra-personal and inter-personal – by which information is collected into and transmitted between controlling conceptions (files) belonging to one or more subjects. For example, I might collect information through perception and put it in a file, and then tell someone about this encounter, which will allow them to put this information in a file of their own. Or I might collect information in a file, and then realize that it is from the same object for which I have a pre-existing file, and merge the two.

Evans opens by discussing three processes of what he calls the information system: through perception we collect information, through memory we retain information, and through testimony information can be transmitted.

In the first part of the section, which runs to the middle of p. 124, Evans distinguishes his notion of information from sensation and belief, by noting that a “traditional epistemologist” would have described these in terms of *sensation* or *belief*. He starts with a discussion of the two ways of *acquiring* information: perception and testimony (memory is a capacity for retaining information). On the traditional account, perception and testimony would require an inference in order for the subject to end up with a belief that had objective significance. Evans claims that it has been widely accepted that in the case of perception the traditional approach has been abandoned, and the “data” of perception — what is given — should be thought of not in terms of sensation (e.g. empiricist sense data) but rather *seemings*, characterized as “events ... already imbued with apparent objective significance”. What this means exactly is not very well spelled out. I will return to this in a discussion section after Chapter 7, since some of the material developed in that chapter will help here.

The important point is that they are not belief states, and the reason for this is that such states, on their own, aren't even *thoughts* (a belief is a thought judged to be true). Evans wants a way to characterize the content of information states that allows the informational content to be necessary but not sufficient for thought. The reason being that the *content of* the information state will not satisfy Russell's Principle. The information state itself may have features *aside from its content* that could play a role in satisfying Russell's Principle — for instance, one could use it as part of a definite description such as ‘the object that was the causal source of this information’; or the information link to the object might be used to allow the thinker to locate the object in space (see Chapter 6). But the informational content itself does not provide the means for independently satisfying

Russell's Principle. There is quite a bit to say here, but I will put it off until the Discussion section for this chapter.

The second related point, discussed from the middle of p. 124 to bottom of p. 125, is that the content of information states need not concern the specific objects that produced the information. Evans' example is the content of a photograph. Note that Evans is not here saying that the content of perception or testimony is exactly like that of a photograph in all respects. He is here making one specific point that informational content need not require specification in terms of the specific object that caused the content. Evans provides one characterization of the content of a photograph, e.g. a photo of a red ball resting on a yellow cube:

1. Red(x) & Ball(x) & Yellow(y) & Cube(y) & On(x , y)

Call this photograph photo1. The content is in terms of predicates (red), sortals (ball) and variables (x). The role of the variables is just to show how information is organized: the fact that x is used for Red(x) and Ball(x) indicates that it is the *same thing* that is depicted as being red and a ball. Crucially though, the *content* does *not* concern any *particular* objects. A complete specification of the *content*, while it would have to mention the properties red and yellow, and sortals such as ball etc., and be organized via variables, would not need to mention any specific objects, including whatever objects (if any) were the causal source of the information. I could take a different photo (photo2), where the camera was pointed at a different red ball on a different yellow cube. And it could be the case that the photos would carry the same content (captured in terms of (1) above), and perhaps even be visually indistinguishable.

In such a case, the content of the information carried by photo1 is the same as the content of the information carried by photo2. Even so, Evans claims, it is an important *fact about the information* (though not part of the *content of the information*) carried by photo1 that it was derived from Spherel and Cubel. Evans suggests that we can capture this aspect of information when we specify that the information is *of* Spherel. This isn't a large departure from normal usage: we could say that photo1 and photo2 are 'the same' (they look the same, have the same content) even though they are *of* different things.

Speaking of a mechanism (like a camera, or our perceptual system) that registers information, Evans says

... we can say that the product of such a mechanism is *of* the objects that were the input to the mechanism when the product was produced. Correspondingly, the output is *of* those objects with which we have to compare it in order to judge the accuracy of the mechanism at the time the output was produced. Notice that I have explained the sense in which a photograph is *of* an object, or objects, without presupposing that a specification of its content must make reference to that object, or those objects. (p. 125, underline emphasis added.)

It is worth emphasizing that in using the word "accuracy" Evans is clearly referring to the accuracy *of the mechanism*, not the accuracy *of the content* produced by the mechanism. [Failure to see this difference is behind a recent misplaced criticism of Evans' by Burge (2010, p. 184-5). Burge interprets Evans as claiming that we assess the accuracy *of the content* by comparing it to the objects. From this, together with Evans' remark that the content does not concern the objects that caused it, Burge finds an inconsistency. But if this

difference is kept in mind, Burge's criticism evaporates. For a fuller discussion, see Grush 2018, 'Understanding Evans.']

That was the first point, that the *specification of the content* of a piece of information need not involve particular objects, including the objects the information derives from. The second part of the section, from the bottom of p. 125 to the end, concerns what Evans calls the "social informational system". This includes mechanisms for collecting, storing, and communicating information within a single agent and between multiple agents. Note that there are elements of the "social information system" that are not social. Individuals have the ability to gather information through their senses and retain information in memory.

There are several interesting points here. Recall from Section 5.1 that Evans describes as "pure cases" situations in which the content of a file (controlling conception) is exhausted by the current content-acquiring episode. This contrasts with what he is now describing (the middle of p. 126) as controlling conceptions that are *mixed*. These are cases in which the content of the file is the result of distinct information-acquiring episodes, and the files have been merged. For example, one has a file for an object that one is perceiving, and after collecting a certain amount of information, one re-identifies it as an object that was encountered on a previous occasion. In this case, the new file that was created when the current encounter began will be merged with the file that was created in the previous encounter. Of course it is possible that the subject will fail to realize that it is the same object, and will maintain two separate files. Note also that files can be split, if the subject learns that what she took to be one object is in fact two — I meet one of two twins on Monday and the other on Tuesday, and I initially take them to be the same person I will merge the files. When I learn I met two people, I can

try to separate the information back into two files — e.g. the one I met on Monday plays basketball, and the one I met on Tuesday had a cold.

Evans claims that singular terms require re-identifications of the (c), (d) and (e) sorts. Note though that the requirement doesn't go the other way, since one can make re-identifications of these sorts without singular terms. I might decide that two files should be merged — that they are files for the same object — because of sufficient overlap in content, or some other sort of inference. For example, I see a big blue bird in the park on Monday, and on Tuesday a friend tells me they saw a big blue bird chasing a dog in the park, and perhaps because birds of that sort are rare, I merge the files and take it that the bird I saw on Monday was chasing dogs on Tuesday. What singular terms do is, so to speak, invoke a dossier directly, independently of the content in the dossier. If these sorts of connections weren't possible, there would be no point to referring expressions.

Evans goes on to argue that in the cases of social information passing, the distinction discussed earlier between information being from *a*, and information having content whose specification requires mention of *a* specifically, still holds — even in cases where it is a particular that is the source of the information, the content of passed information can fail to be particular-involving in that way. In the cases Evans discusses here, the content is existential. Evans sums up the argument nicely:

Consider my belief that swans are to be found in Uganda: is this about some particular birds, long since deceased, that occasioned the report in the travel book from which I derived the belief? [p. 128]

Finally, there are a couple of clarifications about information states. First, the object an information state is *of* (in the sense discussed above) is

determined not by what object, if any, the information *fits*. It is the object, if any, that was the causal source of the information. Also, two different people may have the same bit of information. Here the individuation is in terms of information-gathering episodes.

Here is an example. I see a blue bird on Monday, and see it eating from a trash can. I see a blue bird on Tuesday and see it chasing dogs. But suppose that though it was the blue bird I saw, I misperceived it as *green*. And by coincidence, there was a green bird chasing dogs in that park that day. Even if I don't identify them (which I may not, given my misperception on the second day), the pieces of information are *of* the same object — even though the second bit of information fits the unseen bird better, it is still *of* the (blue) bird that I saw. Now suppose that I told you “A green bird was chasing dogs in the park on Tuesday”. You will now have the same bit of information that I gathered on Tuesday, since it derives from the same “informational event.” What this latter expression means will be explored in Chapter 9.

5.3 Interpretation and psychological attributions

This section is very straight-forward. Evans here wants to explain why in many cases, including the putative counter-examples to Russell's Principle (the second steel ball, the polish grocer), we feel inclined to credit the subject with a thought about the object even though the subject has no discriminating knowledge of the object — an ascription that would constitute a violation of Russell's Principle — even when such an ascription should not strictly speaking (according to Evans) be made. The short version is that we find it natural to use 'relational psychological idioms', relating the subject and an object x , — e.g. “S means x ” or “S wants x ” — when interpreting

people's thought and speech if we know that the subject has an overriding goal which would be best satisfied by reference to x. And this is true even if the subject is not really in a position to think about or refer to x. E.g. a student of auto mechanics who says that "A spark is produced in the carburetor", and another student corrects him by saying "He means the cylinder," which is a natural thing to say even if the first student has no idea what a cylinder is. The first student's overriding goal — to say true things about car engines — would be best served if he had said 'cylinder'. Given that people have an overarching goal of thinking of and speaking of the objects from which their information derives, it is natural to attribute to them thoughts that concern those objects even when, strictly speaking, they are unable to have a thought about the object. So Evans claim is that the reason we find it natural to say that the subject is thinking of the first ball is not that they are genuinely capable of thinking of the first ball, but rather because we know that the subject has the over-arching goal of thinking about the object that is the causal source of their information, and the first ball was the causal source.

The second of the steel ball examples was the one where the subject, in the past, sees a steel ball on one day and a different but identical one the next day. But she retains information only about the first one. We are assuming that the subject has no individuating knowledge of the balls, in that the information content she retains from the first ball matches the second ball equally well. According to Russell's Principle, the subject should not be able to think a thought about either of the balls.

This situation is supposed to present Evans (and RP) with a problem. The problem is that it seems natural to say that the subject is in fact thinking a thought about the first ball in this case, even though the content of the information she is basing her thought on matches the second ball equally

well. And hence she is not distinguishing the ball from all others, as RP requires.

(Recall, the subject could always attempt to think of one of the balls by description by invoking the information link itself, “the ball that was the causal source of the information I have”. This would be a way of distinguishing the ball from all others, and so would satisfy RP. But if we are trying to come up with a counter-example to RP, we need to assume the subject is not attempting such a descriptive thought.) So the problem is: The subject is not satisfying RP, yet there is an inclination to say that she was thinking of the first ball. Evans has to address this.

He does this by showing that there are conditions where it is correct, in a sense, to attribute to someone an intentional state with x as its object even when, strictly speaking they can't really have such an intentional state with x as its object. In particular, our use of psychological idioms, such as “S is thinking about x ”, or “S means y ”, can be ways of us registering the fact that, even though the subject isn't really thinking about x or mean y , they have a larger-scale project that would be satisfied if they did say P or refer to x .

Evans gives two examples. The first is the auto mechanics student who says, when questioned by the teacher, that the spark is produced in the carburettor. There is a clear sense in which the student's overall project and goals involve saying true things about the operation of the combustion engine. And what would have satisfied that goal would have been if the student had said “The spark is produced in the cylinder.” Saying that is what would satisfy his overarching intention, and so it is quite natural to say, as the other student did, that he means the cylinder, or he meant to say that the spark is produced in the cylinder. And we can attribute this to him even if he has no idea what a cylinder is. The other example is of the person who wants to register a complaint, and is told that he wants to see Mr. X. The person did

not (in any strict sense) have an intention to see Mr. X. He did not know of Mr. X, and may not even have known that one needed to see anyone to register a complaint — perhaps he thought that a form needed to be filled out. But it turns out that his over-arching goal of registering a complaint requires him to see Mr. X. So it is natural to say 'he wants to see Mr. X'. So our interpretive project of making sense of people's thoughts and utterances involves a sensitivity to their high-level goals, and these can often license an attribution of an intentional relationship between the subject and objects that would satisfy those goals, even when it is clear that the subject is not capable of a state that is directed at that object or proposition.

Next, Evans points out that people have, as one of their high-level goals, the aim of speaking or thinking about the object which is in fact the causal source of the information in the relevant controlling conception. With these two pieces in place, Evans' response to the counter-examples is clear. It is that yes, we do in fact find it very natural to say that the subject is thinking of the first ball (the one from which the information derives) when in fact the subject has no discriminating knowledge of this ball. But that is not because the subject really is thinking about that ball. Rather, it is because the subject has, we can safely assume, the higher-level goal of thinking about whatever is in fact the causal source of his information. And the attribution is a natural way for us to register this fact. However, the subject is not in fact entertaining a thought that is genuinely about first ball, any more than the student had an explicit intention to say 'cylinder' or the complainer actually had an intention to see Mr. X.

In connection with these points, Evans introduces the notion of well-groundedness: a thought-attempt is well-grounded if the means of object identification employed by the thought-attempt manages to identify the object that is the source of the information in the subject's controlling

conception. Here is another way to put it. In attempting to think about x , I have to employ some means of identification of x . This will be whatever it is that allows me to satisfy RP, the resources I bring to bear that allow me to distinguish x from all other objects. So let us suppose that my means of identification singles out what is in fact a . I also have, as an overarching goal, that the object I identify in this way is the one that is the causal source of the information in the controlling conception of that thought. So suppose I see someone at a party, and overhear them talking. The person says "I am really tired." I distinguish the person from all others by their spatial location, since I can clearly see them. It is that person right there that I am thinking about. And if that person is also the causal source of the information I have, then my thought that that person is tired is well grounded. But suppose I made a mistake, and that at the party is a ventriloquist standing off to the side who said 'I am tired.' Here, my thought would not be well grounded. The causal source of the information is something other than the object that my means of identification singled out. Well- and ill-groundedness of thought is the topic of the next section.

It is crucial to note that on Evans view I can have this overarching goal of well-groundedness even if I do not (and maybe I cannot) articulate this goal as a description of the form "the object that is the causal source of the information I have about x ." Of course, sometimes this is possible, but in such a case the thought would become a descriptive thought. The supposition here is that my thought-attempt is not descriptive but demonstrative.

5.4 The Risk of Ill-Groundedness

Following the discussion of last section, Evans claims that it is an overarching goal thinkers have that their thoughts be well-grounded — that is, that the means used to distinguish the object of thought identifies the object which is in fact the causal source of the information upon which the thought is based. Here Evans spells out some different ways that ill-groundedness can occur. It is crucial to note that the example Evans uses involves thoughts that attempt to identify their objects by description. This makes the example easy for present purposes, but runs the risk, for the incautious reader, of giving the impression that Evans thinks that in general the way people distinguish objects of thought is by description, and this isn't the case. He just hasn't discussed any other ways yet, and descriptive identification is familiar, and so makes for a good example at this point.

The topic is information-based thoughts. There are many types, including demonstrative thoughts, recognition-based thoughts, and others. One kind is what we might call an information-based descriptive thought. *Standard definite descriptions* are of the form *the ϕ* . For example, when we speculate that there was a unique person who invented the zip, we can think about *the inventor of the zip*. An information-based descriptive thought would be one where the description is built from information from the object, perhaps along with details of the encounter. For example, if my neighbor sees a big blue bird and tells me about it, I might think about *that big blue bird*, where the components of the description — big, blue, bird — are the result of information received from the object (in this case from my neighbor's perceptual contact with it), and then transmitted to me via testimony. Evans

will mark the difference between a definite description and an information-based descriptive thought linguistically as

Definite description: (i) *The* ϕ is F

Information-based descriptive thought: *That* ϕ is F

Evans is not suggesting that this is an invariable means English speakers actually use to mark the difference. It's just how he will mark the difference. Nor is it to be assumed in these examples that in all cases of information-based thoughts object is being identified demonstratively (which might be a natural way to interpret 'that'). He is just using this convention to mark the difference between information-based and non-information-based in the examples. Note that the predicate F might also be the result of information gathered during the perceptual contact — my neighbor saw it in a tree and it looked like it was building a nest, and so F would be '... is building a nest'. In the examples Evans discuss it is the *identifying* component — the stuff in the description — that is important. And generally speaking, the well-groundedness goal concerns the harmony between the information source and the object identified by *that* information. That the identified object be the source of the information concerning the predicative component is not a factor.

So the setup is that a neighbor tells me about a big blue bird they saw down by a local river. He says 'That big blue bird down by the river is building a nest.' I then essay a thought that *That big blue bird (down by the river) is building a nest*. Note that it is important for these examples that it is *my* beliefs, not my neighbors, that are under investigation. This is because we want to focus on the identifying information in the description, but my neighbor, who saw the bird, saw its spatial location and had a different sort of

identification (demonstrative) of it. Let's walk through the four possibilities Evans lists, plus a fifth well-grounded case.

1. In the first case, there is a source, x , of the identifying information, but there is no unique object identified by the mode of identification, either because there is more than one object that fits the description, or because there is none. One way the latter case could occur is if x were a green rock that looked vaguely bird shaped from one particular angle to my neighbor with poor vision, and due to a trick of lighting appeared blue from that angle. Either way, there is no unique satisfier of the description, no single 'big blue bird'. And either way, the doctrine is that my thought that *that big blue bird is building a nest* is ill-grounded. There is an information source, but the means of identification doesn't single it out.

2. The second case is like the first in that there is an object x that is the source of the information, except that in this case there happens to be something y ($y \neq x$) that fits the description. My neighbor never saw it, he only saw the rock. It was purely accidental that y was there and happened to fit the description. So my thought *that big blue bird is building a nest* would also not well-grounded in this case. There was a causal source, and the means of identification does in fact single out a unique object, but it isn't the same as the causal source.

3. In the third case, there is no object that is the source of information. My neighbor lied, or suffered a hallucination. (Note that a *hallucination*, where there is no causal source, is different from a *mistake* in which a green rock is mistaken for a big blue bird, for in the case of the mistake, there is a causal source of the information.) Even so, there happens to be an object y that fit the description. Again, my thought that *that big blue bird is building a nest* is ill-grounded.

4. In the fourth case, there is no source for the information (it was the result of hallucination or lying), nor is there anything that fits the description. Again, the thought that *that big blue bird is building a nest* would be ill-grounded.

5. (Not enumerated by Evans) The only case which would be well-grounded is one in which there was one big blue bird, and it was the source of the information. Notice that ill- vs. well-groundedness has nothing to do with the truth of the thought. My thought could be well-grounded but false, if the big blue bird was in fact not building a nest.

Evans and PM agree that in 3 and 4 there is no thought. Both Evans and PM take thoughts to be Russellian — if there is no object that is the causal source, then there is no belief/thought. Evans and PM also agree that in 2 the subject is not thinking of y (the bird that coincidentally fits the description). The disagreements are that in 2, Evans thinks the subject isn't thinking about y because the subject isn't thinking a thought at all. But PM takes it that in 2 the subject has a belief about x . The same is true of 1: Evans says no thought, PM attributes a belief about x . (Remember, Evans can maintain in the 1 and 2 that my neighbor is able to think about x , because my neighbor actually saw the bird and so could discriminate it by its location in space. I am relying on the descriptive material *big blue bird* to identify the object.)

The first thing to note is that Evans will claim to have an explanation for why it seems natural to attribute a belief about x to the subject in the 1 and 2 cases, namely, the sorts of considerations brought up in 5.3 (there is a typo in the text, at the middle of p. 134 there is a reference to 5.4, which should be 5.3.).

The crucial issue is whether it is correct to ascribe a genuine belief about x (not in the 5.3 sense) to the subject. Evans argument against this is this:

If it were, then in a case of type (2), or type (1), in which (i) the subject's attempt at an information-based thought which he would express by saying 'That is F' is based on information which is derived from x , but

which constitutes a wholly inaccurate conception of it; and (ii) x is in fact F , we ought to have an equally strong conviction that our subject thought something true. And I do not think we would wish to say that the subject thought something true in this kind of case, because we would be extremely embarrassed if we had to provide an account of what it was that he thought. (VR p. 135)

Here is what this means. Change the example so that the descriptive material that derives from x constitutes an inaccurate conception of x (because of mis-perception, etc.), but the predicate is nevertheless true of x . For example, my neighbor sees what is in fact a large grey rock in the water, but mistakes it for a woman lounging in the river. My neighbor tells me that *that woman in the river is old*. And let us suppose that the rock is in fact old — both for a rock (it is precambrian) and for a human. It seems that the PM theorist must maintain that when I say ‘That woman in the river is old’ I am expressing a true belief about the rock.

Evans' argument is this: in cases where the conception is fairly close — a woman drinking water from a martini glass is mistaken for man drinking a martini — there is little resistance to ascribing a true belief. I hear about this person, and say “that man drinking a martini is a spy”, there is little resistance to saying I have a true belief, about the woman, that she is a spy. So the force applied by the considerations discussed in Section 5.3 makes it natural to ascribe the belief. But once the example is changed to a case where the conception is extremely inaccurate, we can see the problem. Since the rock is the causal source, and being the causal source is, according to the PM theorist, sufficient for being the object of belief, then the PM theorist must say that I have a true belief, about the rock, that it is old, when I say ‘that

woman in the river is old'. If it was the truth of PM that was driving the intuitions, then both cases should be parallel.

The upshot (which will be expanded upon in the Appendix), is that descriptive information-based thoughts are unlike definite descriptions. Because a successful thought of this sort must be well-grounded in order to be successful, Evans is not endorsing anything like a 'description theory' (which would maintain that the object of the thought is the object that uniquely fits the description), even of this sort of descriptive information-based thought.

5.5 Preview

Evans in this section discusses two kinds of object-dependence that a thought might have. First, an information-based thought is object-dependent because all information-based thoughts must be well-grounded, and if there is no object, being well-grounded isn't possible. A second kind of object-dependence a thought might have is that some Ideas of objects require that the object exist in order to be coherent, and so any thought-attempt employing such an Idea will be object dependent for that reason.

Different referring expressions require different kinds of thought to be understood, and so in some cases they are Russellian for one of the above reasons, in some cases for both reasons.

Recall in Chapter 3 Evans discussed two different ways to establish that a singular term was Russellian. Some kinds of singular terms might be such that in order to be understood, the subject must grasp a certain kind of thought, and the grasping of such a thought might be object-dependent, in

that, if there is no object, there is no thought; other singular terms might be Russellian because, though they don't require any particular kind of thought to be understood, the object itself is what provides the 'principle of unity'. In the first approach, but not the second, the singular term is Russellian because the required thought is Russellian. In this section Evans points out there are two ways a thought can be Russellian. We end up with three kinds of case:

1. The Idea is not object-dependent, but it is information-based. Russellian because of well-groundedness requirement only. Information-based descriptive thoughts are an example.
2. The Idea is object dependent, and the thought is information-based. Russellian because of well-groundedness requirement and because Idea is object dependent. Demonstratives and recognition-based thoughts are examples.
3. The Idea is object dependent, but the thought might not be information-based. Here-thoughts and I-thoughts are examples. Russellian only because Idea is object-dependent.

It should be noted that in the Appendix, Evans will claim that in fact information-based descriptive thoughts, here listed under (1), are really (2), the Idea they employ is not actually available if there is no causal source.

5.A Appendix

The appendix tackles some of the same issues as 5.4 and 5.5, but from a slightly different angle. Any thought-attempt will use (or attempt to use) an Idea of its object, which reflects the RP-satisfying mode of identification used to identify the object. This

is what, according to Evans, supplies the content of the thought. Moreover, some thoughts are information-based, in that it is a requirement that the object, if any, identified by the mode of identification be the same as the causal source of the information.

This way of looking at it treats the two factors as essentially distinct. This comes up in the case of the sort of thought that was used as an example in 5.4, information-based thoughts that employ a descriptive identification (that ϕ), since it seems that the Idea is essentially a description, and the 'that' is just registering the fact that the thought is information-based, and hence needs to be well-grounded. But Evans argues here that even in such cases, that the thought is information-based must be part of the content reflected in the Idea.

Recall from Chapter 4 that any thought requires an Idea of the object, and this Idea, if fundamental, serves to distinguish the object of thought from all others. The Idea provides that part of the content of the thought that corresponds to the object. That was Chapter 4. In Chapter 5 the topic was information, and the claim was that some thoughts are information-based. This means that, regardless of the Idea employed by the thought, if the thought is information-based then it must be well-grounded, in that the object distinguished by the mode of identification (provided by the Idea, and which constitutes the content) must be the same as the object that is the causal source of the information.

As far as this description goes, i) the Idea-provided content, and ii) the fact that the thought employing this Idea is information-based, are independent factors. In Chapters 6 and 7, Evans will discuss Ideas that are constitutively tied to certain ways of gaining information. But even in those cases, these two factors are still independent. This is most clear in the cases of

here-Ideas and I-Ideas which, though tied to ways of gaining information do not require any actual information to be coming in through the information-link.

This duality comes up explicitly for the sort of thought that was the example employed in Section 5.4, information-based thoughts that use a descriptive identification (e.g. that blue bird). The discussion there suggested that the Idea (the means of singling out the object) was a description, the same description that would occur in a definite description *the big blue bird*, and the switch from 'the' to 'that' was just registering the fact that the thought was information-based, it was not registering any difference in the mode of identification.

But Evans here suggests that this can't be quite right, that the fact that the thought is information-based is partially constitutive of the Idea. That is, the Idea employed in 'the φ ' and 'that φ ' are different Ideas. Evans expresses this

When, in the description-based case ... we say that the subject believes that that φ was F, we use 'that φ ' in such a way that nothing could be that φ unless it was the source of the retained information on which the thought is based. (p. 139)

Surely the difference between the thought that *the φ was F* and the thought that *that φ was F* is a difference of content, and not a difference which can be wholly shunted off into some sphere of considerations external to the determination of content. [p.140]

It is important, though, that though the fact that a given Idea supports information-based thought is part of the content provided by that Idea, it does not provide that content by adding a specification of the causal link as

sort of a descriptive element: *that blue bird* does not mean the same as the definite description *the blue bird which is the source of my information*.

If the fact that the thought were information-based were worked into the descriptive content, then it would not be possible for two subjects to employ the same mode of identification, because they would be using different descriptions — each description would include a clause about a causal link between “the object and me”, but since the ‘me’ is different in the two cases, so are the descriptions, and hence so are the Senses (the intuitive criterion of difference would surely distinguish them). Evans puts it this way:

If we conceive the content of information-based thoughts in the wrong way, then it can seem impossible for two people to share an information-based particular-thought. One person will be regarded as thinking that the object which is causally responsible in the appropriate way for his possession of such-and-such information is thus and so; and the other will be regarded as thinking that the object which is causally responsible in the appropriate way for his possession of such-and-such information is thus and so ... (VR p. 141-2)

But if the fact that they are information based is a constitutive part of the Idea, then two subjects can still employ the same Idea of the object, meaning they can think of it in the same way. Specifically in the *information-based* way.

One point of clarification. Earlier in the chapter (Section 5.2) Evans argued that it is possible that two subjects might have the same bit of information. Given this, one could be puzzled about what Evans is saying here. If you and I can have the same bit of information, then surely we might both grasp the description “The object responsible for *this bit of information*.” That is, if you and I can have the same bit of information, why would my description need to make reference to *me*, and yours to *you*?

The response is that in order to be the same bit of information, the information needs to derive from the same information gathering event. I see an object, and then tell you about it. In such a case the bit of information derived from the same event — even if in your case that bit was also transmitted to you via testimony. This is not the same as two people each perceiving the object. In that case, there are two information gathering events — my perception of the object, and your perception of it, and hence the bits of information are not, on this account the same.