

Two Concepts of Belief

I

1.1 Introduction

Contemporary analytical philosophy is replete with philosophical distinctions. Such is, for better or for worse, the result of the systematically performing conceptual analysis in a piecemeal way. The results are not always enchanting. Very often, we find that some of those distinctions are not well motivated or that they could be reduced to other previously introduced distinctions. This was, for instance, the case with the so called distinction between the referential and attributive uses of definite descriptions. As Kripke has shown, the distinction turned out to be unmotivated as a semantic distinction, since it could be handled by exploiting already existing semantic resources such as the *de re / de dicto* distinction or the one between the small scope and wide scope of quantifiers.¹ It could also be shown that if a language contains no singular terms, and if Russell's theory of descriptions can be maintained, the distinction between *de re / de dicto* reports and the distinction between small scope / wide scope occurrences of quantifiers amount to the same thing. Philosophical distinctions can, up to a certain point, be useful tools for the analytical philosopher, but she must sooner or later be ready for a more synthetic approach. Otherwise, we run the risk of losing sight of the more general philosophical picture. By always looking at the trees, we run the risk of losing sight of the forest.

Since I belong to the analytic tradition, these critical observations apply to my own philosophical work. And if I am particularly inclined to worry about these particular philosophical habits, it is because I also happen to belong to a tradition that holds that the meaning of a term should be equated with its uses. I believe that words very often have many different uses and that by uncovering them, we can put an end to some philosophical confusions, or can very often show that certain philosophical debates turned out to be misplaced.

There are two different sorts of confusions to be encountered. The first one is that by failing to make a certain distinction, the two opponents may fail to realize that they are in fact talking about two different things. The second one is that they could fail to notice that they are both developing a different systematic philosophical account on the basis of just a few examples. An illustration of the first sort confusion may very well have been, as Quine thought, the debate between the classical and intuitionistic readings of the words which in our language play the role of logical connectives. As Quine suggests, it may very well be that the intuitionists are not really challenging the classical interpretation, but are rather, in a way, “changing the subject”. They are talking about an entirely different use of the same words. An example of the second kind of confusion could be encountered in the now famous debate between the theory/theory of belief ascriptions and the so called simulation approaches. It could be argued that those two approaches rely on two equally acceptable linguistic uses. There is one which systematically exploits third person reports and presupposes that the

rationality of the agent who makes the ascriptions must always be taken into account in an understanding of those ascriptions, and there is the other one which exploits first person reports and emphasizes the subjective perspective of the agent to whom the ascription is being made. But as a matter of fact, the two uses are both important and must both be accounted for within a unified theory.ⁱⁱ

In this paper, I shall try to motivate what I take to be an important philosophical distinction between two uses of psychological verbs. In particular, I shall try to distinguish between what I call the material reading of belief reports and their intentional reading. I shall first begin by contrasting this distinction with already existing ones, and I shall then put it to work by trying to show its fruitfulness for the resolution of philosophical debates that either rest on confusion or on unjustified generalizations based upon particular uses of words.

1.2 The Distinction

Let me first introduce the distinction that I am now alluding to. I would like to establish a semantic distinction between what I call the material and the intentional uses of mental predicates.ⁱⁱⁱ This terminology is not one with which I am entirely happy, but I have failed to find a better one.^{iv} In what follows, I shall restrict my attention to the particular case of belief. I claim that belief attributions are at least of two kinds. They sometimes serve to report contents that can in some way transcend the cognitive or conscious capacities of the agent. In all these instances, the ascriber reports what is “in fact” believed by the agent. It is what she happens to believe whether she knows it or not. I call this the material use of belief. The ascriber can also try to reproduce as faithfully as possible the actual content of the agent’s belief, the very mode of presentation under which she entertains the belief. In this case, it appears that if the agent knew the language in which the ascriber formulates his belief report, she would be willing to accept the report as such. The agent’s intentional beliefs must thus involve a knowledge that one has the belief. I call this the intentional use of belief.^v

The main criterion that enables us to differentiate the two readings is the fact that in material reports, the ascriber does not assume that the agent has first person authority over the attitudes that she is actually entertaining, while in intentional reports, he does make such an assumption.^{vi} This is the essential difference and I need not identify at this point intentional beliefs with beliefs that are accompanied with dispositions to assent. It is true that in some cases at least, first person authority could be manifested in dispositions to assent, but this does not mean that we should equate the two notions. The dispositional account is just the beginning of an explanation of the phenomenon of first person authority. But my essential concern, at this point in my argument, is to establish the existence of the distinction, and not necessarily to defend a particular explanation of the notions involved.

Here are some examples of material belief reports. The ascriber could make a belief report by using into the subordinate clause indexical or demonstrative expressions that reveal her own perspective on things^{vii}. If Pierre believes that Leonard Cohen’s raincoat is blue, an agent could truly utter that Pierre believes that *this* is blue, while pointing at a particular blue raincoat owned by Leonard Cohen even if she did not assume that Pierre knows that this is

Leonard Cohen's raincoat, or knows that this is blue. We could be tempted to treat the belief ascription as true even if Pierre did not have a disposition to assent, for some reason or another, to the claim that this, again pointing at the raincoat, is blue. To take another example, if a student believes that all philosophy teachers in the philosophy department at Université de Montréal are very good, then I could under a material report claim that she believes that *I* am good, even if she does not necessarily know that I am teaching in that department, and does not even think that I am particularly good.

These examples are not meant to suggest that all belief reports containing indexical or demonstrative expressions in the subordinate clause are automatically to be regimented in the category of material reports, for there could be cases where the agent would be disposed to assent to the belief report if she were in the same situation. For instance, in the above example, I could have been tempted to ascribe to Pierre the belief that this is blue simply on the assumption that he would assent to the claim if he were directly put in contact with the raincoat. So the very presence of indexical expressions in the subordinate clause is not in itself a sufficient condition for the existence of a material use of belief reports. It is just that with indexical expressions in our belief ascriptions, it is very often possible to go beyond the actual recognitional cognitive capacities of the agents.

There are other examples of material reports that do not involve the presence of indexical expressions in the subordinate clause. The agent could be ignorant of some of the consequences of her beliefs and still be believing these consequences. For instance, if she believes that *p* and if *p* implies *q*, one could be tempted relative to our own perspective to conclude that she believes that *q*, even if she is not aware that *q*. Of course, we are not forced to infer that she has such a belief for any *q* entailed by *p*, for if she had just one false belief, it would turn out, under the material reading, that she believes everything, since from a false claim one can logically infer every proposition. I am rather alluding to a subclass of logical consequences. These are consequences that I would want to assume or presuppose in my belief ascriptions. If I want to assume that *p* implies *q*, then I might want to conclude that she believes that *q*.

Or consider a slightly different case. If Pierre believes that Hesperus is a planet and Phosphorus is identical with Hesperus, then we could also truly report that, in some sense, he also believes that Phosphorus is a planet. We could come to agree with this without having to endorse any particular theory of proper names. And since we are not committed to a particular view concerning proper names, the example provides an independent evidential basis for the existence of what I have called the material use of belief reports. The example could be treated as an instance of the first sort of examples only if we endorsed a Putnamian theory according to which there is always a hidden indexical element involved in our uses of common nouns or proper names, but we are not even making such an assumption in the present context.

Another sort of examples comes from unconscious beliefs, or from beliefs that, although registered in the brain, are not consciously entertained by the agent. Let us suppose that, on the basis of his desire for Jocasta, Oedipus comes to believe that he is going to marry her, we could want to conclude in the material sense that Oedipus wants to marry his mother,

given that Jocasta happens to be his mother. This is surely something that he would not want to endorse, and he most probably even would want to deny it, but nevertheless he could in some sense believe it. Oedipus does not know it, but as a matter of fact, this is what he believes.

It might once again look as though this example can be assimilated to the previous ones. Since it involves in the subordinate clause an indexical element (“his mother”), this suggests an assimilation to the first class of examples, and since it is based upon the application of the principle of substitutivity of identity, one could also be tempted to assimilate it with the second category of cases. But it is clear that these aspects of the example are totally irrelevant. It is not simply because of the presence of indexicals in the subordinate clause, or because of an application of the substitution of identical expressions, that Oedipus ultimately would be unwilling to subscribe to this terribly embarrassing belief. It is because of the fact that the belief is repressed. As a matter of fact, it is plausible to suppose that one could have drawn the same conclusion by looking directly at Oedipus’ behavior. In so doing, we could have been led to conclude that he has a repressed desire for his mother and that he unconsciously believes that he is going to marry her. And whether or not one believes in the actual existence of such repressed beliefs, the fact remains that a speaker could want to use her belief report in such a way. So unconscious beliefs provide a good illustration of beliefs that one could be tempted to ascribe to someone in what I have called the material mode.

I have by no means exhausted the examples that could be provided as an illustration for the material use of belief reports. We could, for instance, add to that list the case of innate beliefs. When we ascribe innate beliefs to someone, we are very often in effect making a material belief report if we assume that the agent does not necessarily know that she has the belief. Again, the example is not meant to reveal a special sort of belief, but rather a special use of belief reports. It is not required to subscribe to the existence of innate beliefs in order to acknowledge the existence of such a use.

But let me now turn immediately to examples of intentional reports. When we want to ascribe a belief to someone, we may want to use the very same words to make sure that we capture the content of what she actually believes. In this case, the agent speaks the language of the ascriber and is willing to accept the belief ascription as formulated with the very same words. In such a case, we could in effect be making a belief ascription on the basis of the application of a weak disquotational principle.^{viii} Granted that the agent asserted “p” and that the sentence does not contain indexicals, we can conclude that she believes that p. I suppose that the usual constraints must be assumed by the reporter. She must presuppose that the agent is sincere, reflexive and semantically competent. This, I suppose, is the clearest case of an intentional use of the belief predicate.

Another example would be to ascribe to an agent a belief containing a subordinate clause which is treated as “synonymous” with a sentence that the agent would be willing to accept. In such a case, the reporter assumes that she is semantically competent and knows that there exists such a synonymy relation holding between the two sentences. Again, there need

not exist synonymous expressions. The only thing that matters is that the reporter assumes such a synonymy.

But we can use the verb “believe” in the intentional sense even when we use a sentence that belongs to a different language than the one actually used by the agent. Although the agent speaks another language, the reporter can assume that she would be willing to subscribe to the report if she had known the language of the reporter, or if she had access to his particular translation manual. She would have endorsed the claim that is attributed to her just on the basis of her knowing the relevant pieces of linguistic information. It is important to observe that we do not have to restrict the notion of foreign language in the present case to foreign dialects. By “foreign language”, one could mean a foreign idiolect spoken by someone within a given community of speakers who use the words belonging to the same “natural language”. So the distinction I am making does not involve an implicit commitment to a specific conception of language and even less to a community view of language.

Another example is provided by an ascription made in a situation which roughly coincides to the previous case, but it is also assumed that the agent has a certain logical acumen. Let us suppose that we have all the reasons to think that the agent intentionally believes that all men are mortal. And we also have all the reasons to suppose that she intentionally believes that Socrates is a man. And then finally, we find good reasons to ascribe to her an intentional belief to the effect that the modus ponens holds. On that basis we could be in a position to ascribe to her the intentional belief that Socrates is mortal. This is so because we have all the reasons to believe that she would be willing to accept the belief report. Or we have at least all the reasons to believe that in some way she knows that she has this belief.

In an intentional belief report, the agent may even have failed to utter anything which served as a basis for the belief attribution. In such a circumstance, it is enough if the reporter supposes that the ascribed belief faithfully captures what is actually thought by the agent, even if she did not say so and even if the belief has not occurred to her yet. The essential point is to assume that she knows that she has the belief.

This is not an exhaustive list of examples. I suppose that most of the time, though of course not always, the beliefs that one ascribes to oneself are intentional because they trivially satisfy the criterion of self awareness. But I did not want to produce exhaustive lists. It is also important not to think of these reports as determining once and for all a fixed extension. Beliefs that were once entertained only in the one sense could become after a while entertained in the other sense. The purpose of these examples is simply to illustrate the existence of two radically different ways of reporting beliefs.

2.1 Semantic and Ontological Distinctions

I now turn to a comparative study in which I shall attempt to contrast the distinction that was just introduced with other familiar distinctions available in the literature. I hope in this way to provide additional clarifications concerning the distinction itself. The distinction I am making is initially a semantic distinction and not an ontological one. It is a distinction between belief reports as such and not between different sorts of beliefs. It could entail the existence of an ontological distinction between the two only if we were assuming that the proper semantic account is that of truth conditional semantics and if at least some of those two sorts of belief reports were held to be true, but I am not even making such an assumption. As a matter of fact, I would rather be tempted to replace the traditional truth conditional approach with the semantics of assertibility conditions. But I hope that this preference for an entirely different semantic account will not play an important role in what follows. I am granting that there *may* correspond an ontological distinction between two sorts of beliefs, but I am not committed to it. In other words, the present account is perfectly compatible with the postulation of an ontological distinction between two sorts of belief, and compatible with a truth conditional treatment of belief sentences. My only claim is that we need not even assume that such an ontological distinction has to be made. We only have to register different *uses* of the verbs no matter what actually exists out there in the world.

I will try to remain as neutral as possible concerning the ontological side of the discussion, because my main concern in this paper is to show that there is such a semantic distinction. So, for instance, I claim that under a material reading of belief reports, the ascriber does not assume that the agent knows that she has the belief. He describes the situation as one which does not require such knowledge on the part of the agent. This is compatible with the suggestion that, as a matter of fact, a human agent who entertains a belief always does indeed have this particular bit of self-knowledge. And if this is so, the account is compatible with the idea that all the beliefs actually exemplified by human agents are intentional in our sense. I do not agree with this account, but it is a possible ontological claim at this point given that I am only trying to capture different sorts of reports independently of the ontological status of the belief. Accepting the existence of material reports is compatible with the suggestion that, as a matter of fact, all human beliefs should ontologically be construed as intentional beliefs and thus as beliefs which involve self knowledge on the part of the agents. But such an assumption should not be built in the very content of belief sentences. We have to accept that we can use belief sentences without making such an ontological assumption.

If one were to think that from an ontological point of view, all beliefs are essentially intentional, there could be room for a material use of belief reports. Under such a use, the sentence merely serves to throw light upon just that portion of reality which concerns what is in fact believed by the agent and regardless of her self-knowledge. The report cuts across what could be from a metaphysical point of view two essentially related phenomena, i.e. believing something and knowing that one believes it.

It could also be that what I have called the intentional belief report does not describe a separate ontological reality. Under such a unified ontological picture, material reports could serve to describe functional states of the agents, and intentional reports could serve to describe a subspecies of those functional states. In other words, the account I am proposing is

compatible with the naturalistic reduction of intentional states as such, and I am not here trying to reify the distinction into two separate ontological categories. I said that if one were to accept the distinction and accept truth conditional semantics, then the mere fact that we sometimes treat belief sentences as true would inevitably lead us to accept two different sorts of beliefs. But saying that we are committed to the existence of two different sorts of beliefs is compatible with the idea that, from an ontological point of view, intentional reports describe a subspecies of functional states. As a matter of fact, it is also compatible with the materialist idea according to which intentional beliefs should be eliminated from a truly scientific cognitive psychology.

I said that the distinction I wished to introduce was a semantic and not an ontological distinction. It could be added that I wish to remain as neutral as possible not only regarding the existence of two different sorts of beliefs, but also regarding their nature as concrete or abstract entities. We must indeed contrast the two concepts of beliefs with the usual distinction between those that are actually entertained and are occurrent in the brain of the agent, and those that are not registered at all in the brain of the agent. The first ones would be concrete entities while the second ones would be abstract. We could arrive at an approach of belief in terms of concrete mental representations by naturally extending Davidson's paratactic theory of indirect speech. This is perhaps not an attractive model of belief in general. However, I need not commit myself to either accepting or rejecting this account. I wish to remain as neutral as possible on this score as well.

Finally, the concept of a material belief does not even imply a realist account concerning the existence of propositional attitudes. I only want to introduce two *notions*. The first is the one according to which agents are related to contents that can transcend their cognitive capacities, while the other involves a recognition of those contents. In the material use, the contents are those that they "in fact" believe but it does not mean that there is a fact of the matter about beliefs.

The distinction between the semantic and ontological dimension of the discussion provides a first motivation for the terminology I am using. I do not wish to describe the two uses as "functional" and "intentional", because this already leads us into deep ontological waters. Our distinction is to be understood in analogy with Carnap's distinction between the material and formal mode of speech. By distinguishing between material and intentional belief reports, I wish to convey to the reader the reassuring impression that we are at this point simply dealing with conceptual distinctions and nothing else. Again, I will try as much as possible to remain on the semantic sea shore, keeping my feet on the firm ground of linguistic uses, and to remain as agnostic as possible concerning the big metaphysical debate. In particular, I will not even try to answer the question whether the two concepts involved are compatible with the supervenience of belief reports over material belief reports.

2.2 Conscious and Unconscious Beliefs

Material reports can be true if and only if the agent is not aware of the fact that she has the beliefs. They serve to report beliefs that one happens to have "as a matter of fact" whether or not the agent knows it. But I do not restrict material belief reports to those that are

unconscious. One must distinguish between having an unconscious belief and entertaining a belief without being aware that one has it. Unconscious beliefs are just a particular brand of beliefs that could truly be reported with a material use of the belief predicate. For such a material report to be true, the agent must either repress the belief, or simply fail to assent to the report when confronted with it. It could be a repressed belief, but it could also be something that has nothing to do with repression. Her failure to endorse the belief could be explained by an ignorance of the facts involved, by intrinsic limitations in her cognitive apparatus, or by a failure to have access to the perspective of the reporter.

Intentional beliefs, by contrast, can by definition become conscious. This possibility is built in the very meaning of an intentional use of the belief predicate. But they need not be restricted to the beliefs that one is currently entertaining, to the beliefs that are both occurrent and conscious. Intentional beliefs can also include those that would become conscious if she was simply confronted with the assertion that she has the belief in question.

It should thus be clear that our distinction must not be confused with the distinction between unconscious and conscious beliefs. For one thing, our distinction is just between uses and not between sorts of beliefs. But it must be noted in addition that unconscious beliefs provide just an example which would validate a material report, and that intentional reports are not restricted to conscious beliefs, i.e. to those that are occurrent and “present to the mind” of the agent.

Saying this does not imply that there is no close connection between the two dichotomies. It is indeed important to mention that the criterion for intentional beliefs refers to the presence of first person authority and not just to a reflexive capacity on the part of the agent. We are not simply suggesting that the intentional report would be accurate as long as the agent does have the belief and represents herself as having the belief. An intentional belief is not only reflexive since it is also one for which the agent has first person authority, and this can only be because she is explicitly assumed to be in a position to recognize that she has the belief.

What I have said concerning conscious and unconscious beliefs could also be said regarding the distinction between tacit and explicit beliefs. Tacit beliefs concern prelinguistic representations that are registered in the brain of the agent, that are very often understood as innate instead of being acquired, and that are also very often never explicitly formulated by the agent during her lifetime. While explicit beliefs are, by definition, those that an agent will be disposed to accept and to formulate in her own language. They thus presuppose the existence of language and could not exist without it. I take it that these two kinds of beliefs merely illustrate an additional instance of the general distinction that I wanted to establish. If anything they contribute to motivate the existence of two radically different ways of using belief reports. Our distinction provides the semantic resources that are needed in order to talk about these different sorts of beliefs.

2.3 *De dicto* and *De re* Beliefs

It is important not to confuse the material / intentional distinction with the distinction between *de re* and *de dicto* attitudes. This latter distinction is not even a particular instance of the one that I am trying to introduce. There can be *de re* and *de dicto* material belief reports as well as *de re* and *de dicto* intentional belief reports. For example, Oedipus may believe that he is going to marry Jocasta. And since she is his mother, he also entertains *about* his mother the *de re* material belief that he is going to marry her. This provides a good instance of the material *de re* reading.

But we also have examples of *de dicto* material beliefs. For any sentence “p”, if an agent believes that p and “p” implies “q”, then, this could be sometimes used as a justification for saying that, as a matter of fact, she believes “q” and therefore has a *de dicto* material belief. Indeed there are clearly *de dicto* material reports. Someone may have heard that whales are mammals and thus would be willing to subscribe to the claim that whales are mammals. But then she sees what looks like a big fish in the water and exclaims: “This sort of animal is not a mammal!” Imagine that she constantly makes that observation whenever she sees such a “big fish”. Now since she is always as a matter of fact pointing to a whale, we might be justified to make an inductive generalization. She seems to believe, as a matter of fact, that in general whales are not mammals, in the material sense, since she is certainly not ready to accept that she has that belief. Now the subordinate clause in this case is “Whales are not mammals”, and it is a perfectly “eternal sentence” not containing any proper names or indexical expressions whatsoever. So her belief could intuitively be construed as a *de dicto* belief and much less obviously as *de re*. She may intentionally be entertaining the belief that whales are mammals, but through her linguistic behavior, we are justified to conclude on the basis of an inductive reasoning, that as a matter of fact, she also seems to believe otherwise. She does not know it, but she happens to believe that whales are not mammals.

It is fairly common to acknowledge the existence of intentional and *de dicto* reports. If I say that the actual president of the USA is mortal and I am sincere, reflexive and semantically competent, then I suppose one could conclude that I intentionally have the *de dicto* belief that the President of the USA is mortal. But should we want to say that all intentional belief reports are *de dicto*? Imagine that someone sincerely assents to the application of the predicate “is mortal” to an individual in the context, who happens to be the president of the United States. The person sees the President of the USA “coming at a distance”.^{ix} By an application of the disquotational principle, we could conclude that she has an intentional *de re* belief about this individual. Specifically, we could conclude that she believes about him (pointing in the direction of the individual who is coming at a distance) that he is mortal. She could also be said to believe something about the President of the U.S., as long as she is able to identify him as the President of the U.S. So there seems to be after all good reasons to admit intentional *de re* belief reports.

De re reports are in general naturally interpreted in the material sense, while *de dicto* reports have often been assumed to be used in the intentional sense, but if I am right, there seems to be no reasons not to accept also the material *de dicto* and intentional *de re*. So we have good reasons to conclude that our distinction must not be confused with the *de dicto* / *de re* distinction.

This approach shows, by the way, that beliefs *de se* do not differ in any important respect from other *de re* beliefs. It is often thought that while *de re* beliefs allow for substitution of identicals, *de se* beliefs do not. Indeed, beliefs about one self are most of the time construed as intensional. But they are nothing more than a species of intentional *de re* beliefs. Since there are clear cases of beliefs about one self that are material (as, for example, when an agent has a belief about someone in front of her and does not realize that she is looking into a mirror), the important difference lies not in the object of the attitudes (as a distinction between *de re* and *de se*) but rather in distinction between material and intentional *de re* beliefs.

For the reason that I just gave, it would not help to exploit *de se* beliefs in order to try to capture the material / intentional distinction. The idea would have been to suggest that an intentional belief contains an implicit *de se* belief. But as it was pointed out earlier, an agent could entertain reflexive beliefs without being aware that she has the belief, whether she represses it or whether she simply ignores, as in the mirror example, that she actually entertains a belief toward herself. Saying this however is not denying the fact that intentional beliefs should be partly analysed as involving beliefs about oneself. It is just that the notion of *de se* belief cannot by itself account for our distinction.

2.4 Intensional (Opaque) and Extensional (Transparent) Belief Reports

Let me now turn to another distinction. It is one that comes very close to the insight involved in our own distinction. It could be claimed that it is one that holds between extensional (transparent) and intensional (opaque) readings of belief reports. Couldn't it? Well it depends on what is understood by the words "extensional" and "intensional". Very often, these expressions are used in order to convey many different semantical properties of sentential contexts. A sentential context may be said to be extensional because many extensionality principles apply to it. These principles are usually the substitutivity of identicals, existential generalization and truth functionality. A sentential context then becomes intensional as soon as it blocks the application of either one of these principles.

With such a criterion in mind, it can clearly be said that our distinction does not coincide with the one that is supposed to hold between sentences that are perfectly transparent and those that are opaque, for it would force us to accept that if someone happens to have a false belief in the material sense, then she believes everything. This is indeed a consequence that would follow from treating the sentential context as truth functional. If someone falsely believed in the material sense that *p*, it would happen that, since *p* truth functionally implies *q* for any *q*, she believes absolutely everything. In general, if the context of material belief reports were fully truth functional, one could be said to believe all the logical consequences of one's own beliefs. This is far too strong and goes beyond anything we wanted to capture with our material notion of belief. So if we are to use the extensional / intensional distinction in trying to account for the one that holds between material and intentional belief reports, we have to restrict extensionality to the application of substitutivity and treat opacity or intensionality as the failure of such a specific extensionality principle.

But even then, the intensional/ extensional distinction is not the same as the one I want to introduce, because as we have seen, intentional reports tolerate under certain conditions the application of the substitutivity principle. If the words “physician” and “doctor” are synonymous and if I assume that the agent knows that, I could then ascribe to her the intentional belief that she is going to see a physician on the basis of her claiming that she is going to see a doctor. Intentionality is “in my book” compatible with a certain amount of transparency.

So it appears that material reports are not fully transparent and intentional reports not fully opaque in the traditional sense of those words. But there are further reasons to remain unhappy with an assimilation of our distinction to the traditional transparent / opaque distinction. In the Fregean treatment of belief sentences, we could treat sentential belief contexts as perfectly extensional even when they are used in their intentional sense. For Frege, the substitutivity principle applies between words referring to their indirect referent, i.e. their ordinary sense. So this gives us a further motivation for avoiding the assimilation.

In addition, it is a well known fact that some non psychological sentences turn out to be either extensional or intensional. So the distinction goes far beyond a mere classification of psychological locutions. For instance, first order quantified formulas must be described as extensional while modal sentences are to be treated as intensional. So we would not be telling the whole story if we were in effect merely assimilating our distinction with the one that holds between extensional and intensional uses of belief sentences. We are at best in a position to claim that under a material reading, sentences manifest a limited form of transparency and are quasi-extensional, while in the intentional sense, they manifest a limited form of intensionality and are quasi-opaque, but that would not be very helpful. And at any rate, we would be talking about some of the properties of the two uses, and would not fully explain them. So I conclude that we have a further reason to stick with our own terminology.

2.5 Belief and Acceptance

Our distinction should also not be confused with the one that holds between belief and acceptance. The idea would have been to suggest that what I called material beliefs should be understood as beliefs *per se*, while intentional beliefs amount to what is now called acceptance. Let me very briefly indicate in what sense the two distinctions do not coincide.

First the distinction between belief and acceptance may at best only play a role in the explanation of our distinction. If one distinguishes between belief and acceptance, one can explain the phenomenon of first person authority involved in intentional belief reports. First person authority could partly be explained as based upon a disposition to *accept* the belief report. When someone entertains a belief understood in the intentional sense, she knows what she believes, and the reason may partly be explained by the fact that she is under such circumstances disposed to accept the sentence that describes her belief. If acceptance were itself explained as the result of a stipulation understood in terms of the performance of a declarative (as opposed to an assertive) illocutionary act on a first person belief sentence, then it would indeed seem possible to explain why the agent has first person authority over her own intentional beliefs. For if a person finds herself in authority to make some stipulations,

she will have as it were a “privileged access” to the facts that were made true by the stipulations. She will know that what was stipulated is the case. So first person authority can be explained as the result of being disposed to accept in this special sense the belief report.

Even if we were willing to agree with this particular explanation of first-person authority, it would not justify an assimilation of the material / intentional distinction with the distinction between belief and acceptance. Under my account, it is important to maintain the existence of two different sorts of belief reports that purport to describe two different sorts of beliefs. By contrast, the distinction between belief and acceptance seems to suggest that beliefs are of one sort only and that acceptance constitutes an entirely different phenomenon.

If acceptance is a conscious and concrete physical act performed by an agent in reaction to a certain uttered sentence in the context, and if it involves an understanding of the content of the sentence, then it will normally also involve a disposition to accept the *belief sentence* that describes the agent’s relation to that sentence. So if anything, a concrete act of acceptance provides nothing more than a further evidence for the existence of intentional beliefs. When we find ourselves confronted with an agent that sincerely accepts a given statement “p”, we seem to be in a position to ascribe to her the intentional belief that p. But this does not constitute by itself an analysis of what it is to have an intentional belief. So even if I am committed to treat acceptance as an evidence for the existence of an intentional belief, I am not committing myself to any particular analysis.

I suggested that acceptance could be treated as a more primitive notion and useful as *one* important ingredient for the analysis of intentional beliefs. And indeed, I would certainly be willing to follow that particular line of argument. But I do not wish to commit myself to any particular analysis of what I have called intentional belief reports. I only want to justify the existence of the distinction. In particular, I do not wish to commit myself to an explanation of first person authority in terms of disposition to accept or to assent. At this stage in my argument, my only purpose is to rely on a phenomenology of uses and not upon a particular philosophical account which would involve the claim that intentionality is to be explained in terms of language. The distinction that I am making needs to be acknowledged as fundamental for the point of view of a phenomenology of uses, even if it is not basic from a purely analytical point of view.

2.6 Narrow (notional) and External Beliefs

It is also important not to confuse the distinction with the one between narrow and external psychological states. The distinction between narrow and wide (broad) psychological states was introduced by philosophers who wished to distinguish between states that are individualistically individuated, and those that are individuated in relation with the social environment.^x Now of course, I would certainly not wish to endorse these individualistic motivations, but nor do I need to commit myself to their denial in order to establish the distinction. Let it only be said that my distinction is *compatible* with the view that beliefs are all individuated in relation with the social environment, as well as compatible with the

opposite claim, and this suffices to discard an assimilation of the two concepts of beliefs with those that were introduced by individualist philosophers.

In other words, the two concepts of beliefs are neutral towards the debate between individualists and anti-individualists. Those who sympathise with individualist philosophers and who wish to subscribe to the distinction between narrow and wide psychological states could also accept the conceptual distinction I am trying to develop without affecting in any way their main claims. And the same remarks could be made concerning anti-individualists.

But why should we resist the temptation to conflate material reports with wide functional states, and narrow or notional belief states with those that I have called intentional? The reasons are the following. It surely seems to be possible to admit wide psychological states regarding which the agent entertains a relation of first person authority, as it was emphasized by Burge and Davidson.^{xi} And individualists such as Fodor or Chomsky could very well agree with the idea that I am unaware of at least some of my narrow beliefs. In short, I think that material reports can serve to describe narrow and wide functional states, and that intentional reports can be used in describing states that involve broad contents as well as narrow contents. So our distinction is more general and does not coincide with the internalist / externalist views about content.

3

3.1 The Ramification of Sense into Linguistic Meaning and Content

I now wish to explore some of the consequences that follow from making a conceptual distinction between those two kinds of belief reports. But in order to achieve this and arrive at interesting results, I must first digress and move to an entirely different subject. I want to argue for a generalization of something that resembles Kaplan's distinction between linguistic meaning and content. This semantic distinction was used by Kaplan in order to account for the behavior of indexical and demonstratives expressions. But I would want to propose an extension of it to all categorematic expressions of the language.

It is important to say at the outset that the generalization I am thinking about is for various reasons only very roughly inspired by Kaplan. I only mention Kaplan's own distinction because it approximates the one I am thinking about. In particular, extending his distinction to all the expressions of the language does not mean that I would want to assimilate all expressions to words that implicitly contain indexical terms. It does not even mean that all the words in the language require for their interpretation a relativization to a context which includes the speaker, the time, the place and a possible world. The linguistic meaning of a word, in my sense of the term, is indeed what determines a content relatively to a *context* that includes the *speaker*, but the speaker in this case must not be understood as the parameter of the speaker occurring in Kaplan's notion of context. It would perhaps be better to say that the linguistic meaning of a word is something which determines a content relative to the speakers' intentions.

The notion of linguistic meaning is thus context relative in a deeper way than the one that is usually associated with Kaplan. We could even describe it as belonging to a third level of sense. We could indeed say that the linguistic meaning of a word, in my sense of the expression, is something which, relative to a context in which speakers' intentions come into play, determines linguistic meaning understood in Kaplan's sense, and the latter, in turn, determines relatively to a context a certain content. We could also, alternatively, enrich the notion of context in a way similar to David Lewis, so that an additional parameter such as the speaker's intentions be included. I do not have any favoured strategy to recommend, and it does not really matter for the present purposes.

What is important is to be able to generalize a distinction like Kaplan's to all the expressions of the language and to acknowledge the fact that expressions have linguistic meaning in addition to content. Linguistic meaning, in turn, must be understood as a paraphrase that a speaker associates with to a particular word. The speaker's intentions are thus relevant, since among all the plausible candidates, the speaker has chosen a particular paraphrase.

So there are many important differences between the distinction that I am proposing and Kaplan's distinction. I am suggesting a generalization of Kaplan's distinction between linguistic meaning and content, and not restricting it to indexicals. I am treating all categorematic expressions as involving linguistic meaning in addition to content. And I am finally suggesting that linguistic meaning is determined in part by speakers intention. Still, my notion of meaning is similar to Kaplan's not only because it is in some sense context sensitive, but also because it must be distinguished from the notion of content. If I find it so useful to invoke the distinction that was introduced by Kaplan, it is because I think that, like Kaplan's, my distinction is very general. It is true that Kaplan used his own distinction within a particular philosophical framework, and I shall have to be careful in trying to avoid as much as possible committing myself to those particular philosophical views. But his distinction was also formulated in very broad terms, and it is for this reason that it is at the same time so powerful.

For Kaplan, the linguistic meaning of a term is something which determines a content relatively to a context. And the content is what determines a denotation relative to a possible world. In what follows, I shall concentrate on the notion of linguistic meaning. The notion of content for a given expression may be understood as its contribution to the truth conditions of the sentence in which it occurs (denotation or extension) or, if one is a verificationist, as its contribution to the verification procedures. By "content of a sentence", we might understand the truth conditions or the verification procedure, or even, as Kaplan does, the "proposition" expressed by the sentence, which determines specific truth conditions or specific verification procedures. What's important in these different ways of specifying the content of a sentence is only that we must understand it as determining a certain kind of intension, i.e. a function from possible worlds into truth values. In general, the content of an expression must be understood as an intension, that is, as a function of possible world into denotation.

The notion of content will not be explored any further in the remaining part of this essay, and I do not wish to commit myself to a rich ontology of entities in trying to analyse it.

In particular, I do not want to commit myself to the existence of (singular or general) propositions, and this is why, contrary to Kaplan, I suggested different ways of formulating what content amounts to. I do not want to commit myself to the existence of objective truth conditions either, and this is why I suggested that the notion of verification procedures could serve to capture what is meant by the notion of content. I want to remain as neutral as possible concerning inflationary or deflationary views of content, and I shall thus be happy to talk about contents in terms of functions of possible worlds into denotations. The model theoretic way of characterizing content is fairly neutral and should be roughly understood as it is used by Carnap (notwithstanding Carnap's own characterization of possible worlds in terms of state descriptions). This is already an indication that I am moving away from Kaplan's own philosophical views on the subject.

So let us concentrate on the notion of linguistic meaning. In order to clarify this notion, we must first note that a distinction must be made between two different ways of apprehending it. Sometimes Kaplan talks about linguistic meaning in intuitive terms, and sometimes he gives it a precise model theoretic sense, and talks instead about the "character" of an expression. I find it important to distinguish between these two analyses of linguistic meaning, and I shall thus use the expressions "linguistic meaning" and "character" just like Kaplan does.

According to my use of the expression, the linguistic meaning of a word is like a paraphrase that a semantically competent speaker in a given context is sometimes able to produce in order to "fix the reference" of the word. The paraphrase may take the form of a dictionary definition within a dialect. It may instead specify an abstract determinate entity to be found in a Fregean third realm. It may alternatively be a special kind of paraphrase within the speaker's idiolect. Or it may be interpreted as specifying the conceptual role of the expression for the agent. Finally, it may even be interpreted as an innate inscription in the language of thought that is registered in her brain. These important theoretical variations do not matter here. The important point is that it is a paraphrase of the expression.

The character, on the other hand, is an abstract model theoretic entity. It is a function from context to content. Two expressions have the same character if they determine the same content in the same contexts of use. Some expressions have an unstable character, and others have a stable one. An expression has an unstable character if it determines different contents in different contexts of use, while one that has a stable character determines the same contents no matter how the context changes. In general, Kaplan thinks that indexical expressions and demonstratives have an unstable character while he treats proper names as expressions with a stable character.

Kaplan makes this distinction between linguistic meaning and character, but it is essentially introduced in order to talk about the same kind of entity (i.e. linguistic meaning) in two different ways. He does not dwell on some of the other differences that could take place between the two notions, such as the fact that two different linguistic meanings may determine the same character, or the fact that a single expression may express different linguistic meanings in different contexts of use. One reason is presumably that, for him, expressions such as proper names, definite descriptions or common nouns do not express a

linguistic meaning in addition to a content. Indeed the notion of linguistic meaning is introduced by Kaplan essentially as a useful device for the analysis of indexicals and demonstratives.

But if we do extend the notion of linguistic meaning in such a way that it applies across the board to all expressions of the language, the situation becomes entirely different. For instance, two numerals such as “3” and “the square root of 9” may be understood as having different linguistic meanings, but express the same character, since they have the same content in the same contexts. As a matter of fact, they also have a stable character since they determine a unique content in any context. Another illustration of the difference is shown by two speakers who may associate different linguistic meanings (“rational animal”, “featherless biped”) to the same expression (“man”), and still apply the expression to the same entities in the same contexts. If we were to extend the notion of linguistic meaning to all the words of the language, common nouns could be understood as expressions that have a different linguistic meaning in different contexts of use. This is so because, in a context of use, a different speaker may associate a different paraphrase to the expression. But they would have the same character if the speakers were constantly applying those words to the same objects in the same contexts.

Within Kaplan’s semantic theory, indexical expressions and demonstratives are typically expressions with an unstable character. They determine a different denotation in contexts which involve a variation of the spatiotemporal parameters, while proper names, definite descriptions and common nouns have a stable character. But these expressions may also be such that their linguistic meaning differs from one speaker to another. Apart from proper names whose character is in my view unstable contrary to the assumptions made by Kaplan, these expressions may have a stable character, since they determine relative to different contexts of use the same denotations. But at the same time, their linguistic definitions may vary from speaker to speaker.

These facts were already noted by Frege in the passage of “Sense and Nominatum” in which he acknowledges that fluctuations of sense may take place from speaker to speaker. For Frege, the name “Aristotle” may for someone mean “the teacher of Alexander the Great”, and mean “The philosopher who was born in Stagyrus” for someone else. In Kaplan’s terminology, we should want to say that the proper name has a different linguistic meaning in these two different contexts, even if both meanings determine a single character. Indeed, the two speakers could very well always apply the term “Aristotle” to the same individual, and this suffices to establish that their linguistic meanings determine the same character. The name “Aristotle” could determine the same denotations in the same contexts of use for the two speakers as long as the individual who is the teacher of Alexander the Great is in the same contexts of use the philosopher who was born in Stagyrus.

Now I do not entirely agree with Kaplan’s particular view of proper names, for he thinks that proper names have a stable character and no linguistic meaning. I also disagree with the suggestion that words in general have determinate linguistic meanings and determinate contents. In addition, he seems to believe that linguistic meanings and contents have an ontological status in a Platonic world, along with intensions and propositions.

Finally, he is committed to a particular community view of language. All these features of his theory may perhaps explain why he was led to ignore this peculiar feature of expressions according to which linguistic meaning can vary from speaker to speaker. To put it crudely, for any philosopher who simultaneously denies the indeterminacy of meaning and reference, reifies meanings into abstract entities, and embraces a community view of language, there can hardly be any room for a notion of linguistic meaning that can vary from speaker to speaker and that can be associated with all the expressions of the language. Or at the very least, he can hardly find an interest to inquire upon such a phenomenon. Of course, I do not want to suggest that the community view of language is incompatible with this idea that linguistic meaning may fluctuate from speaker to speaker. But I suppose that it is a phenomenon that someone will not find very important, especially if she endorses Platonism and rejects the indeterminacy of meaning and translation.

Frege recognized that the sense of proper names could vary from speaker to speaker. But he endorsed a descriptive theory of proper names, he failed to distinguish linguistic meaning and content (his distinction between sense and denotation does not amount to the same thing), and he reified sense into abstract Platonic entities. So this is not a better account of proper names, and surely not an account that I would want to use as a model in order to characterize the distinction between linguistic meaning and content for all the expressions of the language. If I refer to the work of Frege, it is essentially because he was able to account for the fact that fluctuations of sense from speaker to speaker. I believe that this is an important phenomenon, and this is why I am interested in introducing a notion of linguistic meaning that is partly determined by speakers intentions and that is not to be confused with the character of the very same expression. Two individuals may have different paraphrases that count as the linguistic meanings of the expressions they are using, and these may in turn determine a very same function of context into content, and thus a same character.

I do not wish to discuss the particular status of proper names, for this would lead us too far away from our main topic. But I suppose that we have to acknowledge the fact that a single proper name may refer to different entities in different contexts of use (e.g. the name “Aristotle”, which can be used to refer to the greek philosopher as well as to the famous businessman (Onassis) who married Jacqueline Kennedy) just like indexical expressions, and thus has an unstable character. And I suppose that it could have a minimal linguistic meaning (e.g. ‘ the individual named “Aristotle” ’). I would also accept the idea that there could be other linguistic meanings associated with proper names and that these could vary from speakers to speakers.

As soon as we are ready to consider the possibility that fluctuations can take place at the level of linguistic meaning independently of the variations that take place at the level of the character of the expressions, this opens up a whole new realm of possibilities. We could be in a position to accept the existence of expressions with a stable character but with linguistic meanings that may vary from speakers to speakers. For instance, the natural kind term “water” could mean different things for different communities or for different individuals (“The H₂O substance”, “a colourless and odorless liquid that quenches one’s thirst”, etc.), even if relative to all contexts, it determines the same substance and is therefore a word with a stable character.

According to the account I am now proposing, we should *in general* distinguish between the linguistic meaning and content of an expression. And in order to establish the fruitfulness of the distinction for all the categorematic expressions of the language, I have tried to stress the fact that a distinction could be made between character and linguistic meaning. Non indexical expressions which apparently have a stable character may still have different linguistic meanings in different contexts of use. If we take very seriously the idea of extending Kaplan's distinction to all the words of the language, then we must first consider whether expressions that do have a stable character may also have linguistic meanings, and second, consider whether this linguistic meaning may vary from one speaker to the other.

Now since, as a matter of fact, I happen to believe that words have such linguistic meanings and, moreover, that some that *may* vary from one speaker to the other, I am indeed inclined to apply Kaplan's ramification of sense into meaning and content to all the expressions of the language and not to restrict its application to indexicals and demonstratives. In a way, this is something that Kaplan himself authorized. He thought that proper names were not indexicals, nor demonstratives, and yet he said that we could describe them as having a stable *character*. We are now in a way just moving a step further in that direction, and suggesting that they could also be described as having linguistic meaning. And we are also accepting in general the idea that many expressions in the language that have a stable character may have different linguistic meanings.^{xii}

Personal pronouns may have paraphrases similar to the ones described by Kaplan himself. For instance, the linguistic meaning of the first person pronoun "I" amounts to something like "the utterer of this token", or something like that. The linguistic meaning of demonstratives such as "this" and "that", at least according to one of their uses, is "the object here and now". The basic linguistic meaning of any proper name N should be something like "the individual called 'N' ". The linguistic meanings of common nouns may be given by something like Putnam's stereotypes. The basic linguistic meaning of the truth predicate is given by the disquotational principle. And then, the basic linguistic meaning of a formula ascribing truth to a given sentence is given by considering an assertion of the sentence itself. The linguistic meaning of quantified formulas is sometimes given by a partial enumeration of their disjunction or conjunction, depending on whether the quantifier is existential or universal. Their possibly infinite disjunctions and conjunctions give their content. (This is what can be said at least when they behave as substitutional formulas). And so on.

Once again, let me emphasize that my use of Kaplan's distinction was only meant as an inspiration and a point of departure for the distinction I myself was struggling to develop. Kaplan talks about linguistic meaning and content as relevant only for an analysis of indexicals and demonstratives, he reifies these entities into real abstract objects, and endorses a consumerist semantical approach that looks very much like a community view of language. I, on the other hand, want to remain as neutral as possible towards these issues. So why should I refer to the distinction as inspired by Kaplan? The reason is simply that Kaplan has also presented abstractly his own distinction in a way that I find neutral. Very roughly speaking, I agree with him that the linguistic meaning of a word, whatever it is and no matter how it is to be construed, determines a character, i.e. a function which itself determines,

relative to a context, a certain content, and this content, in turn, determines a function of context into denotation. This general model may be used within many different philosophical frameworks, for different purposes and in all sorts of ways, and it is because of that that we should endorse it. At this very general level, my only difference with Kaplan is that I want to extend the distinction to all categorematic expressions of the language. For what I am about to show, it is not important to embrace a particular philosophical framework. Quite the contrary, it is essential to remain as general and neutral as possible towards different philosophical views.

3.2 A motivation for a Sentential Theory of Propositional Attitudes

So in what follows, I am going to take it for granted that we are in a position to generalize Kaplan's distinction to all the expressions of the language. I do not wish to reify linguistic meanings into abstract entities. More precisely, I wish to remain as neutral as possible concerning the particular status of linguistic meaning as such. As far as I'm concerned, they may very well be construed as abstract entities, or as dictionary definitions belonging to the linguistic conventions of a community, or as paraphrases made within the speaker's idiolect. Of course, it is important for me to admit the possibility that a word's linguistic meaning could vary from one speaker to another, but this does not yet commit me to treat linguistic meanings as entities belonging to the speakers' idiolects. For instance, Frege recognized that there could be fluctuations from speakers to speakers concerning the sense of a word, but he certainly did not consider the idiolect of the speaker as the most basic notion of language with which to operate. And those who want to subscribe to a community view of language could also be in a position to acknowledge the existence of idiolects. So the mere fact that I allow for the possibility that a word's linguistic meaning may vary from one speaker to another does not by itself commit me to an idiolectal view of language. Of course, I am not committing myself to a community view either.

So let us suppose that this is a correct picture. What has all of this got to do with the two concepts of beliefs that were earlier introduced? I wish to claim that the distinction between the two readings is possible because we believe sentences and that sentences have linguistic meaning and content. Here I am making use of Kaplan's distinction between content and linguistic meaning in order to prove that there is a structural ambiguity in any belief report. If we apply the distinction between linguistic meaning and content across the board to all the categorematic expressions of the language, and if we make the hypothesis that belief statements involve relations to sentences, we are then in a position to account for a general ambiguity affecting belief sentences. We can use the belief sentence in order to represent a relation between the agent and a linguistic meaning, or we can use it in order to represent a relation between her and a linguistic content. Granted that belief reports describe relations between agents and sentences, the ramification of sense into linguistic meaning and content yields two different uses of belief sentences. When the relatum is a linguistic meaning, we have an intentional use of the belief sentence, and when the relatum is a content, we have a material use of the sentence.

In the examples that I gave as illustrations for the two sorts of belief reports, the frontier that separates the material and intentional uses can be seen precisely as the one

between a relation to linguistic meaning and a relation to content. The reason why we are tempted to treat the intentional reports as involving first person authority is that we are usually presupposing that the agent is semantically competent and has access to the linguistic meanings of the sentence believed. If agents believe sentences, if an intentional report describes a relation between the agent and a linguistic meaning, and if it is assumed that the agent is semantically competent, then it follows that all intentional beliefs will involve first person authority. Of course, if the agent does not fully master the linguistic meaning of the sentence she believes, it is as though she took the linguistic meaning to be specifiable only by the sentence itself without paraphrase. In this case it is as though the linguistic meaning could be condensed into the sentence itself as mere syntactic shape. At the other end of the spectrum but still within the general category of intentional uses, the reporter may allow for substitutions to take place in accordance with the assumption that the agent is rational and has access to the linguistic meaning of the logical particles. But there are also middle ground examples where the reporter exploits substitutions that presuppose a relation of synonymy, or where he presupposes that the agent would subscribe to the report if she had access to the translation manual (or to the linguistic meanings of the words of her own language).

The same kind of variety may take place within the realm of material reports. If the reporter describes Oedipus as someone who believes that he is going to marry Jocasta, and if the subordinate clause expresses the same content as the one expressed by “He is going to marry his mother”, then the reporter may be entitled, under a material use of the belief report, to ascribe to Oedipus the belief that he is going to marry his mother. If the agent assents to the sentence “Leonard Cohen’s raincoat is blue” and *this* happens to be Leonard Cohen’s raincoat, I can ascribe to her the belief that *this* is blue even if she does not have access to the perceptual experience that I have concerning *this*. To take one last example, if Fido barks at the bottom of the tree and behaves in a way that reveals a functional relation with what looks like a cat in the tree, the reporter may feel entitled to ascribe to Fido the belief that there is a cat in the tree. In short, a material belief report relates the agent with the content of the sentence ascribed. It is because belief sentences relate agents with sentences and because the sense of a sentence ramifies into linguistic meaning and content that we are in a position to ascribe material beliefs. In this case, we are relating the agent with that component of sense that we have called “content”. And if material belief reports do not involve first person authority, it is perhaps because in general we must not assume that the agent has access to the content of the sentence believed.

It thus appears that by making use of a distinction like the one made by Kaplan and applying it to all the expressions of the language, we are in a position to derive the existence of two radically different ways of reporting beliefs. We can explain why there can be a distinction between material and intentional belief reports. Actually, as I said, we need an additional premise. We must accept the suggestion that belief reports relate an agent with a sentence. In short, we must endorse a sentential theory of propositional attitude sentences. And this is the result that I find fascinating. We have, on the one hand, many different reasons to think that there are two radically different ways of using belief reports, and also many different reasons to draw a distinction like Kaplan’s distinction between linguistic meaning and content. Well isn’t there a connection to be established between the two? We are now

discovering that there could be one, but only in so far as we are willing to accept a sentential theory of propositional attitude sentences.

It is in general very difficult to come up with an argument for a particular choice of logical form concerning belief sentences. But since Kaplan's distinction provides an explanation for the distinction between two sorts of belief reports, the material/intentional distinction indirectly serves as an argument for treating belief reports as relations between an agent and a sentence. The argument for a sentential theory thus runs as follows. If we are able with our belief reports to make two different uses that exploit as *relata* the two main ingredients which precisely happen to be relevant in the analysis of the sense of linguistic expressions, then it must be because our belief reports relate agents with linguistic expressions.

3.3 An Explanation of the Intentionality of Belief in Terms of Natural Language

The above argument does not entail that in the real world agent believe sentences. It is an argument concerning belief reports. And as a matter of fact, it is compatible with the claim that we can ascribe beliefs to creatures that have no linguistic ability. As shown in the Fido example, we can make belief reports in the material sense to such creatures. A distinction must therefore be drawn between a sentential theory of belief sentences, and the suggestion that one could not have a belief if one were not a creature endowed with language. Material belief reports enable us to refute a *prima facie* objection made to sentential theories of propositional attitude sentences. It is sometimes pointed out that those who favour such a theory cannot account for ordinary belief ascriptions to animals. But what looks as evidence against sentential theories turns out to be evidence for the existence of material ascriptions of beliefs. As long as Fido behaves in a way that would be in harmony with what it wants only if the sentence "there is a cat in the tree" were true, we can say under a material use of belief reports that Fido believes that there is a cat in the tree. This belief report can in turn be analysed as establishing a relation between Fido and a sentence. And it appears that report is true because Fido appears to be functionnaly related to the content expressed by the sentence.

This shows the fruitfulness of the present approach. We have tried to distinguish between linguistic meaning and content and show that our beliefs understood in the intentional sense could be understood as related to the linguistic meanings of sentences. The suggestion that something like linguistic meaning could be useful for the analysis of belief reports is certainly not new, and has been forcefully argued by many philosophers. What's new is the suggestion that there are many different uses of propositional attitude sentences, and that all these uses can be accomodated within a single sentential theory. By accepting the idea that belief reports relate agents and sentential components, and accepting that a very general distinction between linguistic meaning and content applies accross the board to all meaningful expressions of the language, we are able to justify the existence of belief reports that broadly fall into two distinct categories of use, the material and the intentional. And because of that, we are also able to account for many particular uses that turned out to be only particular instances of this general distinction: the unconscious and conscious, the innate and

the acquired, the quasi extensional and the quasi intensional. We are also able to account for *de re* and *de dicto* material belief reports as well as *de re* and *de dicto* intentional belief reports. We are able to account for narrow and external belief reports in which “believe” is used in the material sense, as well as narrow and external belief reports in which the verb is used in the intentional sense. And we have just seen how the two uses that were introduced enable us to distinguish between belief sentences ascribed to creatures that are not endowed with a linguistic capacity as well as to creatures that have such a capacity.

We are also in a position to derive fairly strong philosophical conclusions. We have just seen that from the material / intentional distinction and the extended version of the one proposed by Kaplan, we were in a position to motivate a sentential theory of propositional attitudes. But we are now with this result able to say more. It appears that if we accept our particular analysis of intentionality which involve first person authority, we can conclude that there can't be intentionality without natural language at least for a limited class of mental states. For if intentionality involves a reflexivity and even self-knowledge, and if the correct account of belief sentences is the one proposed by a sentential theory, then the agent that knows she has a particular belief will know that she is related to a particular sentence. It will thus appear that she cannot entertain an intentional belief without representing herself as being somehow related to a particular bit of natural language.

It is important to note that the thesis according to which intentional beliefs are to be explained in terms of language is not a trivial consequence that follows from accepting a sentential theory of propositional attitudes, for such a theory yields nothing as regards the ontology of belief. It would automatically imply the existence of linguistic belief contents if we were dealing with an objective truth conditional semantics, and assuming that all those uses are irreducible, but we are not making such assumptions at this stage in our argument. It is important to register many different uses of belief reports regardless of their cognitive import. It may very well be, for instance, that a large class of belief reports can only be useful in the context of a folk psychology and useless in the context a scientific cognitive psychology.

So it turns out that we are unable to draw the above ontological conclusion from the mere adoption of a sentential theory. But can we draw instead the conclusion that intentionality presupposes language just by exploiting the idea that it involves first person authority? I do not think that we are in a position to do that. For first person authority could have been explained as involving some kind of self awareness, some kind of epistemic transparency, a privileged access to one's own mind. We need at least both a sentential theory of propositional attitudes *and* an analysis of intentionality in terms which involve first person authority in order to be able to conclude that no one could have an intentional belief without language. More precisely, we are able to conclude that if there is such a thing as intentionality, then we should give an ontological priority of natural language over intentionality in the case of belief.

I just said that the sentential theory and the suggestion that intentionality involves first-person authority are both essential premises in an argument that purports to show that intentionality is to be explained in terms of language for a particular class of mental states.

But actually, this conclusion would not be valid if we were not willing to accept in addition two different kinds of belief reports. It is crucial to accept that some belief reports can be made in order to ascribe attitudes that do not presuppose such a first person authority. Otherwise, belief ascriptions to creatures without language would be misplaced. So in addition to the two premises, we must add that there can be material reports in addition to intentional reports. The thesis that intentional beliefs presuppose language would not be acceptable if we did not allow for a different use of the verb "believe". So our initial distinction between two sorts of belief reports does also play an important role in an argument that tries to establish a strong connection between language and intentionality in the case of beliefs. It is only because we were led to distinguish between two different readings of belief reports that we were in a position to extract a particular use which could plausibly be explained in terms of a relation that the agent entertains with language. There are many different uses of belief reports that do not fall into this category and that are *prima facie* counterexamples to the claim that there cannot be intentionality without language. But as soon as we make a distinction such as the one we tried to establish, we are in a favourable position to argue for that thesis.

It is also important to note that there may only be a terminological debate involved in the question whether intentionality may or may not come by degree. I do not see a major disagreement between the account that I am proposing and the suggestion that there could be prelinguistic intentional belief states. If, by "intentional", one means nothing more than the notion of aboutness, there can then be intentional states which do not involve language. The important point is that if an organism has first person authority over some of her beliefs, and thus knows what it believes, then that organism must be endowed with a capacity for language.

This result is stronger than the usual claim according to which some intentional beliefs require language in order to be entertained by the agent (viz. The belief that low interest rates induce favorable conditions in the country for regional development). For we are now casting a wider net. Even the intentional beliefs that snow is white, that the sky is blue or that grass is green cannot be held by a creature devoid of language.

Are we able to say anything more than that? Are we able to draw conclusions that concern all intentional states? If we agree to treat all intentional states as involving first person authority, does that mean automatically that this self-knowledge should be explained in terms of a particular assent to a sentence? No it doesn't. It is perhaps only in some special cases that intentionality and language meet. It is only in the case of intentional reports about beliefs, judgements, propositional knowledge and "thoughts" that first person authority should be explained as a disposition to assent. We may be tempted to make intentional reports about perceptual states and this means that the agent knows what she is perceiving, but one may know what one is perceiving without being disposed to subscribe to a perceptual statement that reports her perceptual experience. Consider an agent who is now drinking for the first time a bottle of wine. She may know in some sense what she is tasting, in the sense that she is aware of what's going on, and so she may be intentionally tasting something, but she needs not have any disposition to assent to any particular linguistic description. She perhaps need not even be a linguistically endowed creature. She obviously could fail to assent

to the claim that she is now tasting a Californian Cabernet Sauvignon from Napa Valley, but she could even fail to assent to the claim that she is now drinking wine or a liquid. The capacity to articulate one's state in a language is not a precondition for entertaining an intentional experience in general. It may be possible to ascribe some intentional states to a creature which is totally devoid of language abilities. My claim concerning the intimate connection that holds between intentionality and language concerns only a small portion of intentional experiences. It is only in the case of higher order intentional states such as belief, judgements, thoughts and propositional knowledge that we should be able to conclude to the existence of a strong link between intentionality and natural language.

3.4 An Explanation for the Compatibility Between a Tacit Knowledge about the Rules of Language and the Explanation of Intentionality in Terms of Language.

I would like to end this paper by discussing one last issue. Chomsky's semantical program seems to commit him to a highly individualistic account of the mind. It suggests that we have to choose between a Chomskyan approach to the mind and social anti-individualism. The two views are opposite conceptions of the mind. Is this true? Some philosophers have thought that Chomsky's views were at least compatible with a certain form of externalism.^{xiii} Others have been as bold as to suggest that Chomsky's views were even compatible with a Wittgensteinian view of language.^{xiv} There are still others who distinguish between Chomsky's contributions to linguistics and his philosophical views. In light of his recent individualist commitments, I suppose that we have no choice but to conclude that there is a clash between his views and social anti-individualism. But it is less clear whether the same thing should be said about his contributions to syntax.^{xv}

One usual answer is to distinguish between syntax and semantics. Chomsky's theories, understood in terms of contributions to syntax, could be compatible with social anti-individualism, because the latter concern the semantic properties of mental content. This is however not entirely satisfactory, because there may be a close connection between the syntactic and semantic components of a sentence. We perhaps cannot entirely divorce syntax and semantics. Another suggestion is to treat Chomsky's theories as having implications concerning the *capacity for language*. The theory has implications concerning the different conditions of possibility for language, and it affects the way an individual is able to learn a language. But languages, as such, can be accounted for in ways that are compatible with social anti-individualism, i.e. as a conventional system of semantical rules adopted by a given community. Languages can be interpreted as social practices governed by social norms, and they can play a role in the individuation of some of our higher order attitudes. So even if, as Chomsky claims, we have an innate capacity for language which is to be accounted for in individualistic terms, it is still compatible with the suggestion that our mental contents are individuated in accordance with social anti-individualism.

But as long as we remain vague and general, we are not going to be able to move very far in our understanding of this issue. There seems to remain at the very least a tension, if not a contradiction, between Chomsky's views and social anti-individualism, especially if this latter doctrine is characterized as implicitly involving a community view of language.

We must first be very clear about the expression "social anti-individualism". There are important differences between a Davidsonian version and Burge's version, and it is only the latter which clearly seems to imply a community view of language. It is also true that even according to Burge himself, the most basic thought experiments serve to establish externalism and not social anti-individualism as such. And so by social individualism, I mean a view such as the one expressed in "Individualism and the Mental". We must also distinguish between the *thesis* of anti-individualism and some of the *arguments* for the thesis. Very often, social anti-individualists try to establish their main thesis by showing that mental properties do not supervene upon the internal physical states of the individual. But even if such a supervenience were accepted, social anti-individualism could remain true. The mere fact that mental contents supervene upon physical inscriptions registered in the brain of a particular agent does not show that these are individualistically individuated. These physical inscriptions could have been the result of learning and not be innate, and could thus be individuated by reference to the social practices of the community.

Be that as it may, there seems to remain an essential link between the argument put forward by Burge in "Individualism and the Mental" and a community view of language, while it is clear that Chomsky would reject the community view of language. So let us grant that there is an opposition here that cannot be settled without having to choose between the two approaches.^{xvi} But I do not wish to discuss these issues in the present essay. It is important for my account to remain totally neutral between alternative views of language. I claim that in what I have shown so far, I am not committed to choose between a community view of language or Chomskyan semantics. If I mention this debate at this stage of my argument, it is only to put it aside.

But Gilbert Harman has once tried to establish that there were specific incompatibilities to be noticed between the two approaches that do have some implications for the claims that I am now making.^{xvii} There are at least some versions of social anti-individualism that require that there should be some kind of priority of language over mind, while a Chomskyan account requires there being a prelinguistic knowledge about the grammatical rules that constrain the learning of languages. More precisely, if we give priority to language over mind, as suggested by some of the arguments for social anti-individualism, there seems to be strong reasons to conclude that belief contents must be structured linguistically, while the Chomskyan approach requires that we should treat at least some beliefs as pre-linguistic. It is important to be reminded that in "Individualism and the Mental", Burge argues that his own argument for anti-individualism is neutral concerning the nature of mental contents, and so we should perhaps not include him in this debate. But even if one does not endorse social anti-individualism, there already seems to be a tension between a Chomskyan approach and the so-called priority of language over higher order intentional attitudes. Now since I just gave an argument to support the latter thesis, it looks as though I should be concerned with Harman's claims, and this is true even if I have not committed myself to social anti-individualism.

So this is where there seems to be an apparent opposition between two radically opposed views. On the one hand, we have a view which tends to analyse belief contents in terms of natural language, while on the other hand, we have a Chomskyan approach that

postulates the existence of tacit beliefs in order to account for a prelinguistic capacity for language.

But do we really have two opposite views here? It seems to me that we don't. We did commit ourselves to construe belief reports as relations between an agent and a sentence, but it is only in the case of *intentional* beliefs that the agent is represented as entertaining a relation to a linguistic component occurring in her mind. But are we committed to treat all belief contents as linguistically structured entities? It should be clear by now that we are not. *Intentional* belief reports may be understood as describing dispositions to assent to sentences, and they may take the form of explicit beliefs. But this is not necessarily so for all beliefs. Our material / intentional distinction allows us to acknowledge the existence of tacit as well as explicit beliefs. The beliefs that we call intentional involve first person authority and can be made explicit, but it is not so for the beliefs that are described with a material use of belief sentences. With the latter we are able to ascribe tacit beliefs to agents. We are in a position to ascribe to the agent a tacit propositional knowledge that has nothing to do with an intentional propositional knowledge. In that sense, the account according to which languages, understood in terms of social norms, serve to individuate intentional beliefs is compatible with a Chomskyan account which postulates the existence of tacit beliefs about some of the fundamental rules of language.

Michel Seymour

Université de Montréal

Footnotes

ⁱSaul Kripke, “Speaker's Reference and Semantic Reference”, in Peter A. French et al (eds), *Contemporary Perspectives in the Philosophy of Language*, Minneapolis, University of Minnesota Press, 1979, 6-27; see also my “Pretending to refer” in Herman Parret (ed), *Pretending to Communicate*, Berlin, Walter de Gruyter, 1994, 51-70.

ⁱⁱI have attempted this in “A Sentential Theory of Propositional Attitudes”, *Journal of Philosophy*, 1992, Vol. LXXXIX, 4, 181-201.

ⁱⁱⁱ I try to clarify the distinction in a certain number of papers: “A Sentential Theory of Propositional Attitudes”, *op. cit.*, and “Indirect Discourse and Quotation”, *Philosophical Studies*, 74, 1994, 1-38. I do not claim to be entirely original in these matters, but I never came across a formulation of the distinction in the literature that entirely coincided with the distinction I have in mind.

^{iv} It comes from a distinction once introduced by Elizabeth Anscombe between intentional and material objects of our intentional experiences. See “The Intentionality of Sensation: a Grammatical Feature”, in R.J. Butler (ed), *Analytical Philosophy, Second Series*, Oxford, Basil Blackwell, 158-180. The remaining part of this section can be read in way as an attempt to justify the introduction of a new terminology since, as far as I know, all the existing distinctions can be contrasted with the one I am now proposing.

^v It is fashionable nowadays to claim that “intentionality comes by degree” and thus to allow for different instances among which some do not involve any reflexive capacity on the part of the agent who entertains the so called “intentional” attitude, and much less first person authority. To a large extent, this can be only a terminological issue since, by introducing material reports, I am already conceding that an agent could entertain beliefs without knowing that she does.

^{vi} It is interesting to note that by allowing for those two different uses, we implicitly impose strong constraints upon the resolution of the debate between naturalists and anti-naturalists. We are willing to acknowledge that an agent can have beliefs without knowing that she has those beliefs, as it is now often claimed by naturalists, but we are also acknowledging the fact that there are attitudes that do involve the phenomena of first person authority. Any naturalistic account must therefore provide an explanation of this phenomenon, and it cannot so easily be discarded as it is very often done by the same naturalist philosophers. Of course, saying that there is such a phenomenon does not imply that we should endorse the traditional Cartesian explanation in terms of privileged access.

^{vii} Burge “Self-Reference and Translation”, in F. Guenthner and M. Guenthner-Reutter (eds), *Meaning and Translation*, London, Duckworth, 1978, 137-153.

^{viii} Kripke, “A Puzzle About Belief”, in A. Margalit (ed) *Meaning and Use*, Dordrecht, D. Reidel, 1979, 239-283.

^{ix} This is the phrase used by Russell in “On Denoting” to capture a similar kind of examples which allowed him to authorize exceptionally the substitutivity of identicals in belief contexts.

^x See Jerry Fodor, *PsychoSemantics. The Problem of Meaning in the Philosophy of Mind*, Cambridge (Mass.), MIT Press, Bradford Books, 1987, Chap.2; Brian Loar, “Social Content and Psychological Content”, in R. H. Grimm et D. D. Merrill (eds), 1988, 99-110.

^{xi} Tyler Burge, “Individualism and Self-Knowledge”, *Journal of Philosophy*, 85, 649-663; Donald Davidson argues that Burge is unable to account for first person authority, but he thinks that his own externalist account is compatible with the phenomenon. See his “First Person Authority”, *Dialectica*, 38, 1984, 101-111; “Knowing One’s Own Mind”, *Proceedings of the APA*, 60, 1987, 441-458; “Reply to Burge”, *Journal of Philosophy*, 85, 1988, 664-665.

^{xii} It may be remarked in passing that even if we were to allow proper names to have a Kaplan-esque linguistic meaning, it would not necessarily lead to a vindication of the descriptive theory of proper names.

^{xiii} Patricia Kitcher, "Narrow Taxonomy and Wide Functionalism", *Philosophy of Science*, 52, 1985, 78-97.

^{xiv} Samuel Guttenplan, “Concepts et normes: Wittgenstein contre Chomsky?”, in Daniel Andler *et al* (dir.), *Épistémologie et cognition*, Liège, Mardaga, 1992, 221-238.

^{xv} John Lyons, *Chomsky*, London, Fontana Press, third edition, 1991; see the Introduction.

^{xvi} See, for instance, Burge "Wherein is Language Social?", in Alexander George, (dir.), *Reflections on Chomsky*, Oxford, Basil Blackwell, 1992 175-191.

^{xvii} Gilbert Harman, “Psychological Aspects of the Theory of Syntax”, *Journal of Philosophy*, 64, 1967, 75-87.