[*Nous* 30 (December 1996): 434-460]

Singular Thought and the Cartesian Theory of Mind[1]

Kirk A. Ludwig

Department of Philosophy
University of Florida
Gainesville, FL 32611-8545

1. Introduction

In this paper, I defend some elements of a traditional picture of the fundamental nature of representational states, which I will call the 'Cartesian Theory of Mind'. The theory consists of the following five propositions:

(1) Content properties are nonrelational, that is, having a content property does not entail the existence of any contingent object not identical with the thinker or a part of the thinker.[2]

(2) We have noninferential knowledge of our conscious thoughts, that is, for any of our conscious thoughts of the form 'T(p)' where 'T' represents its mode and 'p' its content, we have noninferential knowledge that we have a thought in that mode with that content.[3]

(3) The content of a thought determines its truth conditions.[4]

(4) We (can) refer and know that we (can) refer to contingent particulars.

(5) Content properties are causally relevant to our behavior, nonintentionally and nonrelationally described.

I think all of these propositions strike us as plausible when we begin to think about the nature of thought and content properties. But difficulties arise when we conjoin these propositions with a number of others that also seem very plausible. In the following, I will first develop three problems for the Cartesian Theory by adding to (1)-(5) a number of plausible semantic,

epistemic, and metaphysical theses (in particular, the now widely accepted thesis that proper names, indexicals, and demonstratives are directly referring terms), from which, together with (3), we can derive the negations of (1), (2) and (5). Among the consequences of these theses will be that there are singular thoughts, i.e., thoughts whose contents cannot be characterized completely independently of reference to some particular object. Independently of the above difficulties, admitting singular thoughts raises a difficulty about how (2) could be true, because it is puzzling what knowledge, particularly noninferential knowledge, of the content of a thought that involves a particular object could come to. I will argue, however, that we cannot avoid commitment to singular thought, given (4), and another extremely plausible epistemic thesis, but I will sketch a conservative account of singular thought that allows us to retain (1)-(4), and show how to resolve the three contradictions we obtain with a minimal alteration to our inconsistent set of propositions. The resulting picture, I will argue, is not only consistent with the semantic and epistemic theses that seem to undermine it, but can be deployed usefully in defense of them. In addition, I will defuse the sense of mystery surrounding knowledge of singular thought and show how this also helps to shed light on some puzzles about self-locating thoughts.

2. Problems for the Cartesian Theory

The first problem is generated by a thesis about the semantics of expressions in ordinary languages and some claims about the relation between propositional attitude ascriptions in ordinary language and thought content. We begin with a claim about the semantics of some ordinary language singular terms:

(6) Proper names, indexicals, and demonstratives in natural languages are directly referring terms.

By saying that a singular term is a directly referring term, I mean that it does not determine its referent solely in virtue of a Fregean sense; that is, there is no property associated with the expression type such that it, together with the way the world is, determines its referent

independently of the use of the term. For our purposes, the only relevant feature of Fregean senses (and, by extension, of Fregean thought contents) is that they are types.[5] By this I mean that in fixing the Fregean sense of an expression (or the content of a thought[6]) one does not restrict in any way any contingent feature of the world.[7] For example, if a singular referring term has a Fregean sense, then nothing follows from a specification of its sense about what objects or kinds of objects exist. Thus, the Fregean sense of a singular term would determine a referent by specifying some general condition that a unique individual satisfies, but which in principle a different individual or more than one individual could satisfy.

I add to (6) the following two claims:

(7) Thoughts are attributed using ordinary language propositional attitude verbs (such as 'believes') and the content of a thought so attributed is given by the proposition expressed by the grammatical object of the verb.[8]

(8) There are some true attributions of propositional attitudes in which the grammatical object of the attitude verb contains a proper name, indexical, or demonstrative, which denotes an object distinct from or 'external to' the attribution's subject.

The first problem is that (6)-(8) entail the negation of (1), the claim that content properties are nonrelational. Consider

(L) Ludwig believes that Nixon resigned in disgrace.

By (7), this attributes to me a content property. By (3), this content property must determine my belief's truth conditions and so must determine Nixon as the object the thought is about. By (6), the proposition that gives the content of my thought is not specifiable as a Fregean sense, that is, independently of reference to a particular object. Moreover, since Nixon, a contingent existent, is not identical with me or any part of me, this means that the content of my thought is not specifiable independently of a contingent existent external to me. Thus, if (L) is true, given (3) and (6)-(8), it follows that I have a content property that is a relational property, contrary to (1).

We will summarize this conclusion as follows:

(A) Some thought contents cannot be specified independently of reference to particular external objects. [(3), (6)-(8)]

Call this 'the externality problem'. The conclusion can be restated as the claim that there are some non-Fregean thought contents that directly involve external objects in their specification (that is, there are singular thoughts about external objects).

To generate the second problem we add to these propositions an apparently plausible epistemic claim,

(9) One's knowledge of the existence of any external object must be based in part essentially on knowledge of one's perceptual experiences,[9]

and an apparently innocent observation,

(10) There are some true attributions of conscious propositional attitudes in which the grammatical object of the attitude verb contains a proper name, indexical, or demonstrative, which denotes an object external to the attribution's subject.

(9) and (10), together with (3), (6) and (7), entail that there are some conscious thoughts we do not have noninferential knowledge of, contrary to (2). For if the contents of some of our conscious thoughts essentially involve external objects ((3), (7), (10)), then if we have noninferential knowledge of those thoughts (of their mode and content) (2), we can infer the existence (by deduction) of the objects that they involve without appeal to sensory experience, contrary to (9). Therefore, given (9), and (3), (6), (7) and (10), we can infer the negation of (2):

(B) There are some conscious thoughts of which we cannot have noninferential knowledge. [(3), (6), (7), (9), (10)]

Call this 'the knowledge problem'.

A third difficulty is raised if we add to these propositions the additional plausible claim that

(11) Causal relevance is a relation that holds between pure qualitative properties or event types.

By a 'pure qualitative property or event type' I mean one that can be specified using a predicate whose analysis contains no singular referring terms. For causal relevance is understood ultimately in terms of causal laws, which govern the evolution of systems independently of which particular systems they are or where in space and time they are located, and so must be specifiable independently of reference to any particulars. The difficulty is immediate. If content properties like that attributed in (L) cannot be specified independently of reference to a particular individual (conclusion (A)), then they are not pure qualitative properties. Hence,

(C) Some content properties cannot stand in the causal relevance relation.

Consequently, some content properties cannot be causally relevant to anything, let alone to our behavior, which contradicts (5). Call this 'the causal relevance problem'.

Finally, there is a difficulty that is not so much an inconsistency as a puzzle about how to understand the conjunction of two of our propositions. This is the conjunction of (2) with our conclusion (A) that some of our thought contents are essentially object involving. Waiving the difficulty about having knowledge of external objects independently of sensory experience, there is the question what knowledge of a thought content that involves essentially an object could come to. Call this problem 'the problem of singular thought'.

To retain (1)-(5) one must give up at least one of (6)-(11). I propose to reject (7), which will be sufficient to render the rest of our propositions compatible (for (7) is used in the derivation of the negations of each of (1), (2) and (5)).[10] Most of the difficulties generated above seem to be connected with the conclusion that there are singular thoughts. It might seem that in

giving up (7) we can avoid the difficulties raised for (1), (2), and (5), because we can hold consistently with the rest that there are no singular thoughts. However, while giving up (7) removes the inconsistency, we cannot avoid commitment to singular thoughts altogether. For we are committed to their existence independently in virtue of (4), that we know we refer to contingent particulars, and another plausible assumption. I establish this point in the next section. Following that, I show how admitting singular thoughts can be compatible with (1)-(4) and explore briefly what sorts of non-Fregean content we have to admit. After that, I return to the question of how we must understand (1)-(11), excluding (7), so as avoid inconsistency and resolve the externality and knowledge problems. Next, I show how the position we reach not only does not conflict with the thesis that there are directly referring terms in natural languages, but can help to deflect the standard arguments against it. In the penultimate section, I take up the causal relevance problem, which requires a modification of proposition (5), but one which does not undermine our explanatory practices. I end with some remarks on knowledge of singular thoughts designed to dissipate its apparent mysteriousness and show how the perspective reached helps to solve a puzzle about first person thought exemplified in Lewis's case of the two gods (Lewis 1983, pp. 140-1).

2. Reference and spatio-temporal symmetries[11]

Our commitment to singular thought, independently of endorsing (3) and (6)-(8) above, depends on two extremely plausible assumptions. The first of these assumptions is (4) above, that we know we (can) refer to contingent particulars. The second assumption is that

(12) We do not know whether our universe is spatio-temporally symmetrical.

The universe is spatio-temporally symmetrical iff every spatio-temporal location t within it is such that there is at least one other spatio-temporal location t' which t cannot be distinguished from purely qualitatively (including all qualitatively specifiable relational properties that do not presuppose the existence of particular objects or space-time points). An example of a spatio-

temporally symmetrical universe would be one for which Nietzsche's hypothesis of eternal recurrence were true, that is, a universe of infinite temporal extent in which every period of time is repeated qualitatively identically an infinite number of times.

If the universe is spatio-temporally symmetrical, then we cannot refer to any contingent particular via a Fregean sense or content. For this would require that we refer to a contingent particular as the unique object that satisfies some Fregean sense or content. It would satisfy a Fregean sense or content by virtue of its general features. However, by hypothesis, every contingent object that exists in a spatio-temporally symmetrical universe is qualitatively identical with another contingent object. Hence, no unique spatio-temporal particular could be determined by a Fregean sense or content in such a world.

If we could refer to contingent particulars only by Fregean content, then we could know that we refer to contingent particulars only if we knew that the universe is not spatio-temporally symmetrical. We do not know that the universe is not spatio-temporally symmetrical. We do know that we refer to contingent particulars. (For example, each of us knows that he refers to himself.) Therefore, it is not the case that we can refer to contingent particulars only by Fregean content. It follows that we can refer to some contingent particulars directly. From this it follows that there are singular thoughts.[12]

Thus, if the basic difficulty is generated merely by the existence of singular thoughts, rejecting proposition (6) or proposition (7) above will not be sufficient to resolve it. We have an argument for the existence of singular thoughts that is independent of reflection on the semantics of our language and the various forms of sentence we use to attribute psychological attitudes.

3. A conservative resolution: Cartesian singular thought

If we are to avoid inconsistency by rejecting proposition (7), while retaining the Cartesian Theory, then in light of the results of the last section, we must explain how this is possible compatibly with there being singular thoughts, i.e., thoughts which are about particular objects but do not pick them out solely in virtue of the types under which they fall. That is, we need to explain how reference to unique spatio-temporal particulars is possible which meets the

constraint that it not be via a Fregean content and that it not presuppose the existence of anything except the thinker. As a design problem, this admits of only one solution, namely, that a thinker is able to refer directly to himself, or to some feature or part of himself, and then to other spatio-temporal particulars by way of a unique relation they bear to him or to that feature or part of himself to which he directly refers. A thought which refers directly to a given spatio-temporal particular, and to another by way of a description of its relation to the first, presupposes for its own existence only the existence of the first particular, but not the second.[13] The second is secured as an object of the thought by satisfying a uniquely denoting description, which makes the same contribution to the satisfaction conditions of the thought whether or not anything satisfies it.[14] Thus, if all unique reference to spatio-temporal particulars is secured by way of direct reference only to the self or a part or feature of it, then no thought of an individual about a spatio-temporal particular will require the existence of any object in addition to him, thus securing that content properties are nonrelational properties.[15]

What kinds of non-Fregean thought must we admit in order to understand how we can refer to or think about the sorts of things we suppose we can? In particular, what about oneself is one able to directly refer to, and what other, if any, non-Fregean features of thought content are required in order potentially to refer to other spatio-temporal objects?

We have noninferential knowledge of propositions expressed by sentences of the form:

(*) I am F

as used on particular occasions, where 'am F' is replaced by a predicate which characterizes a conscious thought. Two things follow from the form of the sentences we use to express what we have direct knowledge of. First, the subject term of the sentence denotes the thinker, and we may therefore infer that we are able to refer, and to know that we refer, directly to ourselves, rather than a part or feature of ourselves. In referring to states of ourselves, or events involving us, we would refer to them derivatively as, e.g., 'my current conscious state', or 'my current visual experience'. Second, we can infer that we can know we have thoughts directly about the present time. For contents expressed by sentences of the form (*) are about the present time; but we do

not know any qualitative description of the present time which we know is sufficient to uniquely denote it, since we do not know that we are not in a universe for which Nietzsche's hypothesis is true. And we could not know we refer to the present time by referring directly to another time and to the present time by description. For even if one could refer directly to a time in the past, say, and to a different time by way of a unique relation the second bore to the first, e.g., as occurring 30 minutes later, one would not thereby know that the time one referred to by description was the present time, since there is no feature of the description which determines that it is.

Therefore, what anchors our references to contingent particulars is our ability to think 'present tense' thoughts directly about ourselves. These are thoughts that cannot be characterized independently of direct reference to the thinker whose thoughts they are and to the time at which they are thought. We can think and talk about events and objects at other times and places by describing them as bearing a unique relation to an event involving us or feature of us at the present time.

Reference to spatio-temporal particulars requires locating them in space, which, in general, requires being able to establish a coordinate system. This requires being able to refer uniquely to a space-time point or region (I drop the second disjunct below). Reference to a temporal point is direct, but reference to a spatial location cannot be, if we suppose that if we refer directly to something in thought we can know that we do so, for there is nothing incoherent in the thought that we could exist, with exactly those states we have noninferential knowledge of, in a nonspatial world.[16] If we knew that we referred to spatial points directly, then this would not be possible. Hence, given our assumption, we cannot refer to a spatial point directly. Unique reference to a spatial location is most plausibly secured by way of our experiential representation of space and of objects as located within it. It is plausible to say that this representation of space is primitive and not grounded in description. Reference to spatial locations can then be secured by way of reference to representations of spatial locations in one's experiential representation of space. This picks out a spatial point, then, only if the spatial experience is veridical.

Reference to spatial and temporal locations requires the ability to determine both a spatial and temporal direction. For we can imagine a universe which is spatially and temporally

symmetrical about our location. We could still refer to spatio-temporal particulars, but only in virtue of being able to describe such particulars as lying in a particular spatio-temporal direction. We will need to be able to represent or think about future and past times, and front/back, left/right, and up/down. These will be non-Fregean elements of the representation. To see this, consider a simple pair of worlds, W1 and W2, which have two spatial dimensions and one temporal dimension, as illustrated in figure 1.

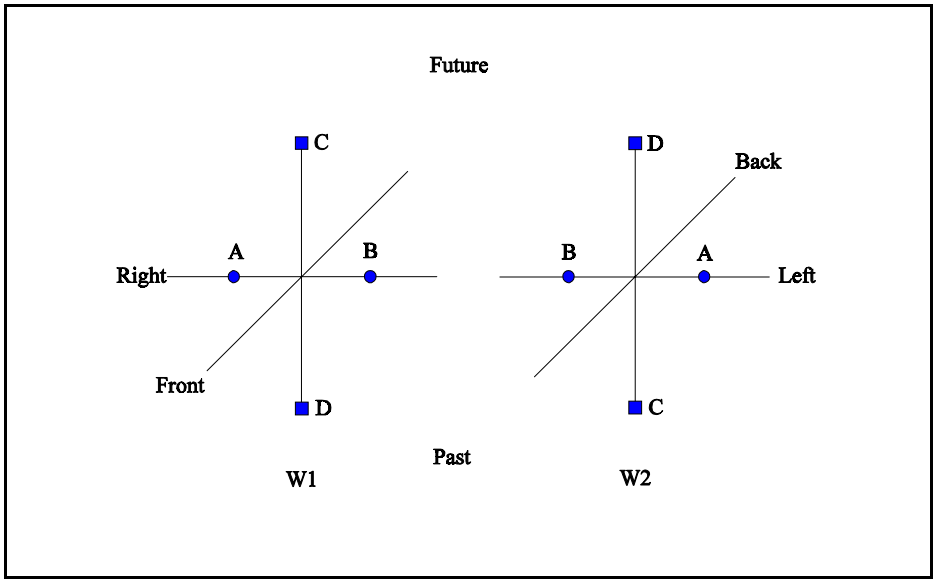

Figure 1

Imagine that you are located in W1 at the intersection of the coordinate lines facing in the direction labeled 'front'. W1 is spatially and temporally symmetrical about the intersection of the coordinate lines (you are qualitatively symmetrical as well). Nonetheless, you would still be able to refer to the object to your right, i.e., to A, and think thoughts about it. Suppose you correctly think: the object now to my right is F. Here, your thought is about A. W2 is by hypothesis qualitatively identical to W1. The difference between the two is that A is located in W2 where B is in W1 and vice versa. Since Fregean contents are types, the Fregean elements of your thoughts in W1 and W2 are the same. However, in W2, you would be thinking about B, rather than A, with a thought of the same type. Therefore, your reference to A in W1, and to B in W2, is not secured by a Fregean thought content. Since reference to the self and a time does not provide sufficient grounds to refer to A rather than B, one's ability to think of a particular direction

represents an additional non-Fregean element of the thought content. While our universe is not locally symmetrical in the way illustrated above, it seems clear that even if it were, we would be able to refer to objects located to the left rather than to the right. Therefore, these points apply to us as we actually are. The same argument can be run with respect to thoughts about the future as opposed to the past. This shows that not all non-Fregean content is so in virtue of presupposing the existence of one or more objects. This ability to think of directions in space and time is best understood as a primitive feature of our experiential representation of space and time.

4. Psychological reality and propositional attitudes

The preceding considerations establish that we are committed to non-Fregean thought contents, if we believe that we know that we refer to spatio-temporal particulars, but they also show that this is no obstacle to understanding how thought contents could be nonrelational. In this section, I explain the basis for rejecting proposition (7), that the contents of our thoughts are given by the propositions expressed by the grammatical objects of ordinary attitude verbs in true attitude attributions.

The Cartesian Theory is about the fundamental nature of psychological reality. By this I mean that it is a characterization of those psychological modes of representing the world that are explanatorily basic in understanding ourselves as rational agents, and in understanding our epistemic position with respect to the world around us. In rejecting (7), I reject the claim that our ordinary attitude verbs correctly track thoughts in this sense. To mark this, let any psychological attitudes which satisfy propositions (1)-(3) be called 'Cartesian Thoughts'. (For convenience below, I will sometimes use 'Thought', capitalized, as an abbreviation. I will continue to use 'thought' beginning with a lower case 't' as a general term covering all kinds of psychological state.) To reject (7) while retaining (1)-(3) is to hold that our philosophical theory of mind should be couched in terms of Cartesian Thoughts. A defense of the Cartesian Theory must do four things. First, it must show that there are Cartesian Thoughts. Second, it must show that they are explanatorily prior to psychological attitudes attributed using ordinary attitude verbs. Third, it must explain why the semantics of our attitude verbs do not track Cartesian Thoughts.

Fourth, it must explain the relation between attitudes attributed using ordinary attitude verbs and Cartesian Thoughts. I take these up in turn.

(i) Why should we accept that there are Cartesian Thoughts? There are two reasons. The first reason is that it is a conceptual possibility that the world seem to us exactly as it does, and yet none of our beliefs about external objects be true because no objects of the appropriate sorts exist. In such a world, we would have no thoughts whose contents have to be characterized by reference to particular external objects. But we do not imagine that we would have no thoughts at all or that our thoughts would not have truth conditions. How the world seems to us is, by hypothesis, the same in both the actual world and in the imagined one, and we imagine having, and knowing that we have, very many of the same thoughts we actually have. Our knowledge of such thoughts would be noninferential, and, in particular, clearly not dependent on knowledge of any external objects. Such thoughts would be Cartesian Thoughts, since they meet requirements (1)-(3) in our initial list: their contents are nonrelational, we have noninferential knowledge of them, and their contents determine their truth conditions. The second reason is that Cartesian Thoughts play a special role in our understanding of ourselves and others as rational agents (and therefore also as linguistic beings). The expression of rational agency in action or thought depends upon attitudes being coordinated for inquiry and action on the basis of their contents. Their coordination for inquiry and action is an expression of our noninferential access to their contents in conscious thought. Attitudes whose contents are not immediately accessible in conscious thought are not suitable for understanding ourselves as rational agents. Since we are rational agents, it follows that we have Cartesian Thoughts. This establishes the first point.

(ii) Why are Cartesian Thoughts explanatorily prior to thoughts attributed using ordinary attitude verbs? The most fundamental reason is that they are, as just noted, the thoughts which figure in our deliberations about what to do. Thus, in understanding ourselves as rational agents, and in understanding our speech behavior, we must look to Cartesian Thoughts. Second, as our first thought experiment shows, Cartesian Thoughts play a special epistemic role in our explaining our thoughts about the world around us, i.e., they form part of the evidential foundation for our beliefs about the external world. Thus, Cartesian Thoughts are explanatorily prior to attitudes attributed using ordinary attitude verbs, since the latter cannot play these special

roles.

(iii) However, if, not only are there thought contents in the sense I have articulated, but they are fundamental to our understanding of our knowledge of our environment and to understanding the nature of rational agency, why is it that our primary ways of talking about each other's mental lives license attributing states some of the contents of which we do not have noninferential knowledge? The answer is that it is the natural outcome of the truth of the account, the pressures on language exerted by the need for efficient and pertinent communication, and our relatively poor epistemic position with respect to others' Cartesian Thoughts.

Suppose that we are a community of Cartesian thinkers set the task of designing a language to enable us to communicate with one another. Given that different people have different life histories and so come to know different things about the same objects, in cases where what is important is what is being said about an object, as opposed to how an individual picks it out, it would be useful to introduce linguistic devices for referring which function in the language simply to denote the object. There are three sorts of cases. In the first, we introduce an expression for denoting a particular which we need to refer to on a number of different occasions. The function of such terms is indistinguishable from that of proper names. In the second, we introduce terms for adventitious reference, where we need not anticipate referring to the same particular on a number of different occasions, but which function, as it were, as proper names for the moment. The denotation of such a device used on a particular occasion is determined entirely by the speaker's referential intentions. His auditor determines the referent of such a device as used on a particular occasion by trying to figure out what the speaker intends to direct his attention to. Such devices function like the demonstratives 'this' and 'that' in English. The third case involves the introduction of terms for referring to objects relative to contexts in systematic ways. Thus, e.g., it is clearly useful to have a way to refer systematically to oneself when one is speaking, although one cannot always presume that one's audience knows one's name. A device introduced for use in the language for this purpose functions like the first person pronoun in English. Clearly, there is a similar need for rule-governed context relative referring terms to fill the various other roles which indexicals, such as those represented by tense, and by 'here', 'now',

'you', 'yesterday', etc., actually play in natural languages.

Having introduced such expressions to talk about the world around us, we will want to introduce expressions to talk about what others think. In our language for talking about the world around us, sentences as used on particular occasions are truth value bearers, as are our thoughts. It will be natural to press into service to describe our thoughts about the world the language we have introduced for talking about it. As in the case of talk about objects around us, in talking about our thoughts or the thoughts of others about those objects, what may be of most importance is what we think about an object, rather than how we refer to it. Thus, it is natural and convenient to press into service in describing each others' thoughts sentences using directly referring terms. When we use such devices, we only partially specify the content of each other's Cartesian Thoughts. But this serves most of our purposes well enough. In addition, often we don't have any very good idea of how an individual picks out an object, though we know what object it is he is thinking of and what he wants to say about it. This is what happens when someone uses a proper name or other directly referring term to pick out an object to say something about it. We know which object, but because the semantics of the referring term underdetermines how the user picks it out, we will often not know how he is picking it out on a particular occasion. In such cases, our best evidence for what the speaker thinks is what he says, and we do no better than to use his sentence (or a sentence synonymous with it) to characterize his mental state, even if in doing so we undercharacterize his Cartesian Thoughts.

Therefore, if we start with the Cartesian Theory, we would expect to see just such directly referring devices in natural languages as we do, and to see them used after propositional attitude verbs. So these facts do not constitute evidence against the Cartesian Theory, and can be explained by it.

(iv) The strategy I have followed in responding to the initial problem has been ecumenical. I have preserved (1)-(4) while denying that our ordinary talk about beliefs, desires, and so on, is in any way flawed. Instead, I have treated 'thought' and 'thought content' as they appear in (1)-(3) as technical terms, introduced to keep track of a special sort of content which plays an important role in our understanding our epistemic relations to the world and ourselves as rational agents. Since I have been using 'Thought' to cover psychological attitudes in any mode,

to discuss the relations between Cartesian Thoughts and ordinary attitudes, it will be useful to introduce technical terms for picking out Cartesian Thoughts in the modes corresponding to belief, desire, etc.; let these be 'Cartesian Belief', 'Cartesian Desire', and so on. On the present view, corresponding predicates formed from the schemas 'believes that p' and 'Believes that p' (the latter representing Cartesian Belief) are not coextensive. What then is the relation between them? How can we describe Cartesian Thought contents in ordinary language terms?

Cartesian Thought contents and ordinary attitude contents are the same sorts of things: therefore, every Thought content can be expressed using a sentence in a natural language. Consequently, to take Cartesian Belief as an example, any Cartesian Belief an individual has is also a belief he has. However, not every belief he has is a Cartesian Belief that he has. Thus, the Cartesian Beliefs an individual has are a subset of the beliefs that he has. They are those beliefs of whose contents he has or can have noninferential knowledge. Thus, many of the beliefs that are attributed to us give also the contents of our Cartesian Beliefs, as, e.g., in the case of beliefs about mathematical truths, and in the case of some beliefs attributed to us (correctly) using directly referring terms that refer to us in the content clauses. The difference comes out only when beliefs are attributed using terms that directly refer to objects not identical to the believer, or to times other than the time of the thought and to places. In these cases the belief contents are not also Cartesian Belief contents. Mutatis mutandis for other psychological attitude modes.

To sum up our progress so far, I have shown (a) that there are non-Fregean thoughts, (b) that this does not entail that thought content is relational, (c) that there are Cartesian Thoughts, i.e., thoughts meeting conditions (1)-(3), and that these are explanatorily basic in our understanding our epistemic relation to the world and ourselves as rational agents, and thus (d) that rejecting (7) is both justified and sufficient to render consistent (1) with (3), (6) and (8), and at the same time (2) with (3), and (6), (8) and (9), solving the externality and knowledge problems.[17]

5. Attitude attributions and entailment relations

On the account given above, directly referring terms can be intersubstituted in attitude contexts

salva veritate. It is often objected that this is a mistake. In this section, I sketch a strategy for responding to the linguistic evidence advanced for this claim. This has a two-fold purpose. First, it helps motivate the solution I offer to the problems generated at the outset. Second, with the Cartesian Theory as background, in the light of the rejection of (7), we can offer a more persuasive account of the linguistic evidence than is usually advanced by direct reference theorists. Given that there is independent evidence for the direct reference theory, this constitutes indirect confirmation of the Cartesian Theory.

I consider two sorts of putative failures of substitutivity.[18] The first involves different proper names. The second involves proper names and indexicals in cases in which the subject of a sentence attributing an attitude and the subject of the sentence giving the attitude content are the same.

The first kind of case can be illustrated using the following two sentences:

(13) Lois Lane believes that Clark Kent is a mild-mannered reporter.
(14) Lois Lane believes that Superman is a mild-mannered reporter.

We imagine a world in which the familiar Superman stories are by and large true, and in particular a world in which Lois Lane is unaware that the mild-mannered reporter she works with is the man of steel she is secretly in love with. It is often said that it is obvious that (13) does not entail (14), and that it therefore follows that either the surface form of these sentences is misleading, or that proper names aren't directly referring terms (see, e.g., (Crimmins 1992), (Plantinga 1974)). The second kind of case can be illustrated using the following two sentences:

(15) Ludwig believes that Ludwig is about to be eaten by a great white shark.
(16) Ludwig believes that he is about to be eaten by a great white shark.

It is said that (16) does not follow from (15) (even evaluated relative to the same time), and, in particular, that (15) cannot be sufficient for (16) because the behavior we would expect on the basis of the attitude attributed in (16) is not in all cases the same as that we would expect on the

basis of the attitude attributed in (15) (see (Perry 1979) and (Richard 1983)). If Ludwig does not know that he is Ludwig, then he will go calmly to his death despite having the belief attributed in (15), while he will not be so fortunate if he has the attitude attributed in (16).

In both cases, we can explain why we treat utterances of the sentences in these pairs as having different truth conditions as the result of differing standard conversational implicata, in Grice's sense (1989b), which we would expect, given that they are in fact semantically equivalent.

(i) If my account is correct, then attitudes attributed using proper names in the content clause (other than the name of the person to whom the attitude is attributed) never fully specify the content of a Cartesian Thought. In particular, such an attribution does not show how its subject picks out the object denoted by the proper name. But as far as the semantics of the attribution goes, it will not matter which proper names we use in the content sentence as long as they codenote. Now, given that different people will have different ways of picking out the referent of a proper name and will typically associate different sorts of information with the individual they think is so named, we can see that someone may announce something about one person using two proper names which she does not realize have the same referent and even be led by this to say things which are necessarily false. It is thus easy to see how someone could rationally say:

(17) Clark Kent will never be as famous as Superman.

The explanation is that the way the speaker has of picking out Clark Kent which she associates with 'Clark Kent' is different from that she has of picking out Clark Kent which she associates with 'Superman'. The most salient associated Cartesian Thought someone has when uttering this is best represented with two descriptions in the place of 'Clark Kent' and 'Superman'. In choosing this sentence to announce the belief, the speaker is exploiting standard descriptions she Thinks associated in her audience's mind with the names 'Clark Kent' and 'Superman' to implicate something not explicitly said by the sentence interpreted literally, e.g., that the reporter working at the Daily Planet known as 'Clark Kent' will never be as famous as the superhero lately of

Krypton, whose amazing exploits everyone is familiar with, and who is known as 'Superman'.[19,20]

When reporting others' attitudes using proper names, we want to give our audience as much guidance about what sort of behavior to expect from them as we can. For this purpose, it is useful to attribute to them attitudes using the sentences which they use in announcing them.[21] This gives our audience some guidance about how the person to whom the attitude is ascribed will react when receiving information about an individual, such as Clark Kent, when he or she hears that information stated using 'Clark Kent'. Thus, there will arise a standard conversational implicatum to the effect that a person who is a member of our linguistic community to whom we attribute an attitude using a sentence such as 'Clark Kent can leap tall buildings in a single bound' would announce her belief using that sentence, and so associates that bit of information with the ways she picks out the referent of 'Clark Kent'. Thus, when we deny that (14) follows from (13), we react the standard conversational implicata of (13) and (14), and not to the truth conditions of (14) and (13) literally understood. (13) and (14) are equivalent, but their standard conversational implicata are not.[22]

If part of the import of uttering a sentence on a particular occasion is an implicatum, then according to Grice it can be cancelled explicitly. To test whether our hypothesis is correct, therefore, we can see whether it makes sense to cancel explicitly the implicatum of (14) and thereby render it unproblematic. I believe we can do this. If someone were to say to you, 'Lois Lane believes that Superman is a mild-mannered reporter' without any qualification, then you would take him to have said something false. However, if he said, 'Lois Lane believes that Superman is a mild-mannered reporter, but she doesn't associate being a mild-mannered reporter with the ways she has of picking out the referent of "Superman"', the false implicatum is canceled, and we no longer feel that what was said is false; instead, it draws our attention to a piece of dramatic irony. This point is reinforced by noticing that there are circumstances in which clearly no such implicatum is present, even without explicit cancellation—for example, when reporting the beliefs of someone who has seen Clark Kent and believes he is remarkably handsome, but who has encountered neither of the names 'Clark Kent' or 'Superman'. In this case, we (who are in the know) could report his belief to one another indifferently using either

name: So and so believes that Superman/Clark Kent is remarkably handsome.

Let us turn now to the second, more interesting case. I need to explain why we think the attitude attributed using (15) has different behavioral implications than the attitude attributed using (16). The difference here is palpable. How is this compatible with the claim that (15) and (16) are necessarily equivalent? The answer lies in the pragmatics of utterances of (15) and (16). But to see how this difference arises, we must take a detour through a considerations of first person announcements of beliefs. Consider a first person announcement of a belief such as,

(18) I believe that I am about to be eaten by a great white shark.

On my view, if I uttered (18) (at t), my utterance would have the same truth conditions as an utterance (at t) by me of,

(19) I believe that Ludwig is about to be eaten by a great white shark.

Nonetheless, the import of my sincerely uttering these two sentences differ. The reason has to do not with the conditions under which the two sentences are true, but with the conditions necessary for me to utter sincerely each of them. In the case of (19), because the semantic role of proper names is simply to contribute a referent, I can be in a position to utter (19) correctly even if I pick out the referent of 'Ludwig' by description. That is, even if 'Ludwig' denotes me, I may know what the referent of 'Ludwig' is only by way of a description of an object that bears a unique relation to me as I refer to myself directly. Since the Thoughts which enter into my deliberations about action are Thoughts which directly refer to me, I can be in a position to announce (19) truly without the thought that I am about to be eaten by a great white shark entering into my deliberations about what to do. If another person reports the belief I announce in (19) as in (15), the same point holds: one cannot from the report infer that I have a Thought whose content attributes to me, referred to directly in the Thought, the property of being about to be eaten by a great white shark. However, the case is different when we turn to (18). For it is clear that anyone who utters sincerely a sentence whose grammatical subject is the first person pronoun is

expressing a Thought in whose content he refers to himself directly. This is so because speaking is a form of rational action, and, as we concluded above, thoughts which are suitable for guiding rational deliberation about action are thoughts of which we have noninferential knowledge and thoughts in which we refer directly to ourselves, that is, Cartesian Thoughts. Someone who understands 'I' knows that it is used to refer to the person who uses it. Someone, therefore, who uses it in a speech act uses it to refer to himself as Thought of directly as being the agent of the speech act. Thus, we can conclude that a speaker who utters (18) has a Thought in which he refers to himself directly, and so a Thought which will play an appropriate action directing role for him. This will be true even though this fact is not represented in the semantics of the first person pronoun 'I'. No such inference is justifiable in the case of (19). That is why we can infer different things about what an agent will do on the basis of his utterance of (18) and of (19), although they are necessarily equivalent.

Now return to (16). This is the form in which we would typically report someone's belief if he expressed it as in (19). On my account, it is equally well reported using (15). If this is correct, then if someone uses one of these forms rather than the other, it will be natural to look for a reason for his using the one rather than the other, something he wants to convey by that particular choice of a sentence to attribute the belief. What could guide the choice of (16) over (15), if my account of their semantics is correct? In (15), despite the use of the same proper name both as the subject of the sentence as a whole and of the embedded sentence, there is no guarantee that the same person is referred to, for different individuals may be designated by the same name. In the case of (16), however, the pronoun 'he' is used as a variable of cross reference to secure that the subject of the embedded sentence is identical with the subject of the sentence as a whole. The point of using (16) then is to indicate that the person to whom the belief is attributed would report his belief using a sentence in which the embedded sentence employed a pronoun of cross reference. That, I have argued above, is just the condition in which we would be able to infer that the speaker was announcing a belief in which there was a direct reference to him. In using (16), therefore, the reporter implicates that the speaker would report the belief using the first person pronoun as the subject of the whole sentence and of the embedded sentence, and that fact carries information about the way in which he refers to himself in the

Thought which he so expresses.

Let us return to the point of uttering (15) instead of (16). We said generally that one will want to report another's beliefs using the words he would use to express them in making statements about the world. Thus, if someone chose to report my belief using (15), he would standardly be implicating that I would announce my belief that way as well. If I did announce my belief in this way, it would be natural to suppose that I did not know that I, referred to directly in my Thoughts, was the person I knew to be named 'Ludwig', as I picked him out. Thus, if I would report my belief by saying 'I believe I am about to be eaten by a great white shark', another's reporting that belief knowingly using (15) instead of (16) would be doubly misleading, since it would implicate both that I would announce it using 'Ludwig' as the subject of the content sentence and that I did not know that I (referred to directly) was named 'Ludwig'. Thus, we see why we resist the inference from what is said in uttering (15) to what is said in uttering (16). The standard implicata are different.

To test this hypothesis, as before, we should see whether this implicatum is cancelable. I believe it is. If someone said to you, 'Ludwig believes that he is about to be eaten by a great white shark' in circumstances in which you knew that Ludwig would not sincerely utter 'I believe that I am about to be eaten by a great white shark', then you would take him to have said something false. However, if he said, 'Ludwig believes that he is about to be eaten by a great white shark, but would not sincerely assert "I am about to be eaten by a great white shark"' the implicatum is cancelled, and we no longer feel that what was have said is false. In this latter case, you would most naturally assume that Ludwig believed of himself in some indirect way that he was about to be eaten by a great white shark.

Thus, the linguistic evidence is compatible with the hypothesis that indexicals, demonstratives, and proper names are directly referring terms and are everywhere intersubstitutable salva veritate provided that they are coreferring. Furthermore, given that there is compelling evidence for treating such expressions as directly referring terms, the fact that the Cartesian Theory provides a natural account of apparent counterexamples to it provides further support for the Cartesian Theory and the relation between it and ordinary attitude attributions that I have sketched.[23]

6. Singular thought and causal relevance

We turn now to the causal relevance problem. This problem remains for us even after rejecting (7), because we have concluded that we are independently committed to the existence of Cartesian singular Thought. This, together with our assumption (11), that the causal relevance relation holds only between pure qualitative properties or event types, and (3), that content properties determine truth conditions, entails (C) that some content properties of Thoughts are not causally relevant to our behavior.

While the conclusion is unavoidable, it poses a serious problem for causally explaining our behavior by reference to our thoughts and their contents only if the content involves the existence of objects other than oneself. To see why, consider the thought content expressed by,

(20) Venus is the second major planet from the Sun.

Someone would be in a state with a content given by the proposition expressed by this sentence only if Venus and the Sun existed. This is clearly not a purely qualitative property, and not object independent. Thus, it is not a property that can figure in the causal relevance relation. The same difficulty arises for the thought attributed with,

(21) I believe that I am presently awake.

For clearly the sentence giving the content contains a directly referring term, which is essential to fixing the attributed Thought's truth conditions, and hence the content cannot be specified independently of an object. However, in this case, the object is identical with the subject to whom the thought is being attributed. The form of the attribution is

(22) x believes x is F.

Although this is represented as having two argument places, in the special case of beliefs about

oneself the second argument place is redundant, and this can be treated as a one-place predicate. A sentence of this form attributing a thought to someone would entail the existence of the individual to whom the property is attributed, but of no other individual. This is in effect to attribute a nonrelational, purely qualitative property to that individual, and a property which is not object dependent. Since Cartesian Thoughts which are about spatio-temporal particulars will all be expressible using sentences that share this feature with (22), it follows that we can construe all attributions of Cartesian Thoughts as attributions of nonrelational properties. Adding direct reference to the present time poses no difficulty, since causal laws quantify over times, and so predicates that figure in them have an argument place for time. The only remaining non-Fregean element in thoughts we have identified is the directionality of certain thoughts, but this presents no difficulty since their directionality does not introduce any direct reference to objects.

While we must give up (5), we can replace it with a slightly modified principle which preserves everything we need in order to make sense of our practices in giving causal explanations for our movements by appeal to our Thoughts. Call properties expressed by sentences of the form (22) (or by sentences in which the content clause contains only predicates expressing purely qualitative properties) content* properties. Our new principle is:

(5^*) Content* properties are causally relevant to our behavior, nonintentionally and nonrelationally described.

Given the Cartesian Theory, it is easy to see that we can construe our explanations of people's behavior by appeal to their attitudes as ultimately resting on the attributions of content* properties to them, which does not violate the metaphysical principle that causal relevance holds between purely qualitative properties or even types. This would not be possible, however, if the properties attributed had to be construed as relational properties. This constitutes, then, our solution to the problem of causal relevance.

7. Knowledge of singular thoughts

The final problem, the problem of singular thought, is not an inconsistency among the propositions I wish to hold, but rather a puzzle about knowledge of thought content when it involves particular objects. The puzzle is not (only) a puzzle about how we could have noninferential knowledge of contents that involve external objects, but also about what noninferential knowledge of an object involving thought could come to, whether or not the object is external. In part the puzzle is generated by a certain picture of what is involved in knowledge of thought contents. Different thoughts can share the same elements. If we think of the shared elements of different thoughts as concepts, then knowledge of the content of a thought involves mastery of the concepts involved in it together with a grasp of the way they are combined in the thought. Concepts of necessity potentially apply to more than one particular. For a thought which determines a subject by means of a description involving concepts, such as 'the first chancellor of the German Empire', this model of knowledge of a thought content will account for our ability to think of particulars. But in the case of a thought which does not determine its referent by some object uniquely fitting a general specification, how is reference to the subject and knowledge of the thought content to be understood?

What I have to say about this is negative, in the sense that I will be urging that various reasons to think there is something mysterious about this are not good reasons.

One approach to this puzzle would be to offer a reductive account of the knowledge one has of a proposition involving oneself in terms of the special role that such attitudes have in deliberation and action. The suggestion would be that all it means to say that someone refers directly to himself is that certain thoughts of his play a special role in guiding his behavior and deliberations, in the sense that they are coordinated in rational patterns with many of his other attitudes. However, while this is a consequence of the noninferential knowledge we have of thoughts which directly refer to ourselves, I do not believe that it exhausts our understanding of it. One thing that is lost is the ability to explain the coordination between some of our beliefs and desires, etc., in producing action as the effect of our noninferential knowledge of their contents. For if that knowledge consists in that coordination, it could not explain it.

If we cannot explain away noninferential knowledge of thought contents in which there is direct reference to oneself, the best we can do is to minimize its mysteriousness and to explain what it is not.

It is not what Russell meant by 'acquaintance'. As Russell put it, "I say that I am acquainted with an object when I have a direct cognitive relation to that object, i.e. when I am directly aware of the object itself" (Russell 1917, p. 152). For Russell, acquaintance was a relation of direct knowledge between an individual and an object. Russell thought (plausibly) that if someone could judge a certain content (a proposition), then he had to be acquainted with the constituents of the content. If we consider a content involving just general concepts, then it is natural to say that one cannot judge that content without understanding or knowing the concepts involved in it. If we then admit into the constituents of a content particulars, it seems, by parity of reasoning, that we should require direct knowledge of those particulars. Reasoning like this must have led Russell to his conclusion that there were some contingent particulars (ourselves, or sense data) of which we had knowledge by acquaintance.

A simpler, less mysterious account is possible. For when we turn to ask what our knowledge or understanding of concepts consists in, the answer seems to be that it consists in turn in propositional knowledge of a special type, namely, a priori, noninferential knowledge of the truth of a set of propositions (or some set in a range of sets of propositions) constitutive of having the concept in question. Such knowledge of concepts is not to be understood on analogy with perception of particulars. The same thing can be said about knowledge of the constituents of thoughts when those constituents are particulars. That knowledge simply consists in immediate (though not now a priori) knowledge of certain propositional contents, contents which themselves contain those particulars as constituents (or require reference to them in their specification), just as we spelled out understanding of concepts in terms of contents which contain those concepts. Acquaintance with the self, to the extent that it is sensible to talk this way, does not consist, as on Russell's view, in a kind of direct nonpropositional awareness of the self, which underlies our ability to think contents involving the self directly, but simply in noninferential knowledge that the self exists, and perhaps noninferential knowledge of certain other contents involving the self, where the self is directly referred to in those contents.

The appearance of a mystery is generated in two stages. First, we take knowledge of contents involving only concepts as our model, and then are puzzled about where to fit in contents involving particulars. Second, when we generalize our model for concepts, we fail to generalize it fully, thereby introducing needlessly a relation of acquaintance with bare particulars. That is, (i) we take our knowledge of concepts as constituents of thoughts as our model for our knowledge of particulars as constituents of thoughts, but (ii) fail to carry through on the analogy by cashing out knowledge of particulars in terms of knowledge of propositions involving them. However, in both cases knowledge of the constituents of thoughts can be understood in terms of certain sorts of systematic knowledge of contents involving those constituents. If we admit, as we are forced to do, contents of both sorts, we should not hesitate to spell out knowledge of contents of both sorts in parallel fashion.

In the end, the best argument for the intelligibility of knowledge of contents involving objects directly is that we are committed to it. We must take such knowledge as primitive, and explain other things in terms of it, but not seek to understand it in terms of any more basic or fundamental epistemic states.

David Lewis (1983) has suggested that knowledge of the sort we express using the first person pronoun constitutes a special sort that is not propositional. If what I have been saying is correct, this is a mistake. It will be instructive to see how it arises. Here is Lewis's case:

> Consider the case of the two gods. They inhabit a certain possible world, and they know exactly which world it is. Therefore they know every proposition that is true at their world. Insofar as knowledge is a propositional attitude, they are omniscient. Still I can imagine them to suffer ignorance: neither one knows which of the two he is. They are not exactly alike. One lives on top of the tallest mountain and throws down manna; the other lives on the top of the coldest mountain and throws down thunderbolts. Neither one knows whether he lives on the tallest mountain or on the coldest mountain; nor whether he throws manna or thunderbolts. (p. 139)

Lewis concludes that there is a kind of self-locating knowledge each can have, which consists in

knowledge of the self-ascription of a certain locational property.

The mistake arises from failing to recognize that our knowledge (even the knowledge of gods insofar as we understand what such beings could be like) has to be structured in a certain way. We know immediately only the contents of our conscious minds. Reference to spatio-temporal particulars depends on direct reference in thought to oneself at a time. All of one's knowledge of spatio-temporal particulars involves essentially contents which directly refer to oneself. Knowledge then is essentially ego-centric. All that it could mean to say that a god is omniscient is that he knows everything about every object. But the contents of a particular god's Cartesian Thoughts would not be the same as the contents of any other god's Cartesian Thoughts, since each of two gods' Thoughts about the contingent world would involve ego-centric Thoughts involving a different god. Since each would know of the existence and nature of the tallest mountain by way of a description of its relation to himself, each would of necessity know whether he was located on top of that mountain, for each would also know what the relation of everything was to himself. The appearance that the gods could be omniscient and yet not know where they were is generated by a picture of the contents of propositional knowledge as purely objective, either because they are purely descriptive, or because they involve bare objects as constituents, where every thinker potentially can entertain any content. Then it appears that something we know is left out, because this leaves no room for privileged thoughts about the thinker himself. But everything we know is left out, because none of our Cartesian Thoughts, when they are about contingent existents, is purely objective in the way envisaged. They are all perspectival, all essentially ego-centric. So, rightly understood, there is no danger in the case of the two gods for the thesis that all our knowledge of our psychological states, and, in particular, knowledge we express using the first person pronoun, is propositional knowledge.

8. Conclusion

We began with a puzzle, an inconsistent set of propositions all of whose members were plausible. The solution consists in rejecting the naive view of the relation between ordinary language attitude attributions and the philosophical theory of thought and thought content, and, in

particular, in rejecting proposition (7). The basic difficulty is not due to the existence of singular thought, which we are committed to provided that we are committed to (4) and (12). To maintain (1)-(4), we must hold that we refer directly only to ourselves at the present time. This, together with the ability to refer to directions in space and time, is in principle sufficient for referring to other contingent particulars by description. This conservative solution to the problem of singular thought, together with the rejection of (7), removes the difficulties facing both (1) and (2). We saw that we could explain plausibly our ordinary ascription practices in terms of the Cartesian Theory, and that linguistic evidence advanced against proposition (6) is also predicted by the Cartesian Theory. The obstacle to understanding how our causal explanations of our behavior by reference to content properties could be correct is also removed by understanding direct reference as being restricted to oneself and the present time. Finally, there is no good reason to find mysterious the claim that we have direct propositional knowledge of the contents of singular thoughts about ourselves at a time.

References

Bach, K. 1988. Thought and Reference. New York: Oxford University Press.

Burge, T. 1978. Belief De Re. Journal of Philosophy 75: 338-62.

———. 1979. Sinning Against Frege. Philosophical Review 88: 398-432.

Castaneda, H. 1968. On the Logic of Attributions of Self-Knowledge to Others. Journal of Philosophy 65: 439-56.

Crimmins, M. and J. Perry. 1989. The Prince and the Phone Booth: Reporting Puzzling Beliefs. Journal of Philosophy 86: 685-711.

———. 1992. Talk about Beliefs. Cambridge,: MIT Press.

Evans, G. 1982. Varieties of Reference. Oxford: Clarendon Press.

———. 1985. Understanding Demonstratives. Collected Papers, 291-321. Oxford: Oxford University Press.

Frege, G. 1980. On Sense and Reference. In Translations from the Philosophical Writings of Gottlob Frege, ed. P. Geach and M. Black. Oxford: Blackwell.

Grice, P. 1989a. Studies in the Way of Words. Cambridge: Harvard University Press.

———. 1989b. Logic and Conversation. In (Grice 1989a).

Kaplan, D. 1978. Dthat. In Syntax and Semantics, ed. P. Cole, 221-43. New York: Academic Press.

———. 1989. Demonstratives. In Themes from Kaplan, ed. by J. Almog, J. Perry, and S. Wettstein. Oxford: Oxford University Press.

Kripke, S. 1977. A Puzzle About Belief. Meaning and Use, ed. A. Margalit, 239-283. Dordrecht: D. Reidel.

Kripke, S. 1980. Naming and Necessity. Cambridge: Harvard University Press.

Lewis, D. 1983. Attitudes De Dicto and De se. In Philosophical Papers, vol. II, 133-159. Oxford: Oxford University Press.

Loar, B. 1985. Social Content and Psychological Content. Contents of Thoughts, ed. R. Grimm and D. Merrill, 133-159. Tuscon: The University of Arizona Press.

Ludwig, K. 1993. Direct Reference in Thought and Speech. Communication and Cognition 26: 49-76.

McDowell, J. 1986. Singular Thought and the Extent of Inner Space. In Subject, Thought and Context, ed. by John McDowell and Philip Pettit. Oxford: Oxford University Press.

———. 1991. Intentionality De Re. In Searle and his Critics. ed. E. LePore and R. Van Gulick, 216-225. Oxford: Blackwell.

McKay, T. 1981. On Proper Names in Belief Ascriptions. Philosophical Studies, 39: 287-303.

McKinsey, M. 1994. Individuating Beliefs. Philosophical Perspectives 8, Logic and Language: 303-330.

Perry, J. 1977. Frege on Demonstratives. Philosophical Review 86: 474-97.

———. 1979. The Essential Indexical. Nous 13: 3-21.

Plantinga, A. 1974. The Nature of Necessity. Oxford: Oxford University Press.

Richard, M. 1983. Direct Reference and Ascriptions of Belief. Journal of Philosophical Logic 12: 425-52.

———. 1990. Propositional Attitudes. Cambridge: Cambridge University Press.

Russell, B. 1917. Knowledge by Description and Knowledge by Acquaintance. In Mysticism and Logic, 152-167. London: George Allen & Unwin.

———. 1918. Acquaintance. In Logic and Knowledge, ed. R. Marsh. London: George Allen & Unwin.

———. 1950. Inquiries into Meaning and Truth. London: George Allen & Unwin.

Salmon, N. 1986a. Frege's Puzzle. Cambridge: MIT Press.

———. 1986b. Reflexivity. Notre Dame Journal of Formal Logic 27: 401-29.

Schiffer, S. 1977. Naming and Knowing. Midwest Studies in Philosophy 2: 28-41.

———. 1987. Remnants of Meaning. Cambridge: MIT Press.

———. 1992. Belief Ascription. Journal of Philosophy 89.

Searle, J. 1983. Intentionality. Cambridge: Cambridge University Press.

———. 1991. Response: Reference and Intentionality. In Searle and his Critics, ed. E. LePore and R. Van Gulick, 227-241. Oxford: Blackwell.

Soames, S. 1987. Direct Reference, Propositional Attitudes, and Semantic Content. Philosophical

Topics 15: 47-87.

Strawson, P. F. 1959. Individuals, an essay in descriptive metaphysics. London: Methuen.

Notes

1. Parts and versions of this paper have been given at (i) The 3rd Karlovy Vary Symposium, "Reference", June 1994, Karlovy Vary, Czech Republic, (ii) at a Rutgers University philosophy department colloquium, August 1994, (iii) the Pacific Division Meetings of the American Philosophical Association, April 2nd, 1995, and (iv) the Central Division Meetings of the American Philosophical Association, April 28th, 1995. I wish to thank the audiences at these talks for their questions and comments. I would also like to thank the following people for helpful comments and advice at various stages in the development of this paper: Kent Bach, John Biro, Walter Edelberg, Jerry Fodor, Terry Horgan, Barry Loewer, Eugene Mills, and Greg Ray.

2. A content property is attributed to a thinker just in case the thinker is attributed a thought with a content. Thoughts in a given psychological mode are individuated by their contents. The content of a thought attributed to a thinker with a sentence of the form 'x φs that p' is given by the proposition expressed by the sentence that replaces 'p'. (A word of caution: this is not equivalent to proposition (7) below, as will emerge in due course.) I link the term 'content property' with the term 'thought'. Therefore, whether belief sentences attribute content properties depends on whether they attribute thoughts. The bearing of this will become clear in section 4.

3. The expression 'knowledge of thoughts' is to be understood this way throughout.

4. Thus, the content of a thought need not be relativized to a context, place or time to determine a truth value, or, more broadly, satisfaction value (for desires, intentions, etc.) for the thought.

5. It is not part of my aim to argue that this characterization of a Fregean sense represents Frege's view. (There is a tradition of interpretation in Britain in particular that runs counter to this characterization.) Rather, I want a convenient label for a certain view of how the truth conditions of thoughts are fixed which represents them as being fixed by content that would be the same no matter what particular objects existed in the world.

6. Thus, importantly, two individuals could, in principle, have the same Fregean thought content, and an individual could have the same Fregean content in any possible world at any time or place at which he existed.

7. I put it this generally because a term can fail to function in virtue of a Fregean sense even though it does not require the existence of any particulars to have the meaning that it does. See the discussion of thoughts about spatial and temporal directions below (section 3).

8. Talk of 'propositions' is not crucial for the argument. The crucial point is that the meaning of the grammatical object of an attitude verb fixes what content property is attributed using a sentence in which it occurs. Propositions in my use are reifications of (eternal) sentence meanings.

9. This traditional view has been disputed, but, while I believe it is correct, I cannot pause to defend it here. For those for whom it lacks appeal, the paper can be read in part as an exploration of how to respond to problems generated by its acceptance.

10. (McKinsey 1994) and (Loar 1985) each reject (7) as well, though for different reasons and in different ways from each other's and my own. A more popular option has been to reject thesis (6), e.g., (Plantinga 1974). Many others have rejected the Cartesian Theory.

11. The debt to (Strawson 1959) should be clear.

12. The argument strictly requires only that what we know about the universe's qualitative character is not sufficient for us to rule out there being two regions of the universe that are, so far as our knowledge extends, spatio-temporally indistinguishable, even if they would not be for an omniscient being.

13. I wish 'description' to include any kind of intentional content, including elements involving experiential representation which may not be fully expressible in any words or concepts.

14. Leaving aside some inessential doctrines, this is the position that Russell put forward in his (1917) "Knowledge by Acquaintance and Knowledge by Description," in response to basically the same problem. Russell had committed himself to our having noninferential knowledge of the contents of our judgements, and to their contents determining the truth conditions of our judgements. Furthermore, he held the traditional empiricist view that we do not have noninferential knowledge of external things.

15. In addition, suppose, for a *reductio*, that one could not refer directly to oneself, but could refer and know one referred directly to some other object O. Suppose one could refer to oneself as the x that uniquely bears R to O. How could one know that one had thereby referred to oneself? By hypothesis, there is an O' qualitatively identical to O. One knows that one has referred to the x that uniquely bears R to O. But there is an object which is qualitatively identical to the x that uniquely bears R to O, namely, the x which uniquely bears R to O'. Which of these is oneself, if either? Nothing that one knows determines this. Therefore, one can know that one refers uniquely to oneself only if one knows that one refers to oneself directly.

16. I owe this point to Eugene Mills.

17. A parallel strategy can be employed with respect to any general terms whose semantic properties are like directly referring terms in the sense that their extensions are not determined by anything like a Fregean sense.

18. I consider only two sorts of alleged failure, but extending the treatment to other cases is straightforward.

19. Of course, this is not a full specification of a Cartesian Thought, which I forebear from representing.

20. Although anyone who uses a proper name N competently must know for some description D that the referent of N is D, the individual's way of picking out the referent of the name is not semantically associated with the name, and, except in special cases, will not be the way the referent of the name is fixed. Competence is easy to come by: in picking up the use of a name from a group of people, we know that the name refers to what the name designates relative to the use from which we acquired it. Since this description does not fix the name's referent, but only secures competence in its use, Kripke's (1980) strictures against its use for the former purpose do not apply. The considerations advanced here also provide a principled resolution of the puzzle presented in Kripke's (1977) "A Puzzle about Belief."

21. This line of response has been anticipated by Tom McKay (1981), although McKay does not develop it within the framework of Grice's theory of conversational implicature. See also (Salmon 1986a, chapter 8).

22. This is not to say that in reporting someone's beliefs we will always be understood to be implicating that he would report it using the words we do. This is so only if it is assumed or known by one's interlocutor that one is speaking about someone who shares with the speaker the language of the report. In reporting the beliefs of someone who does not speak one's language, what is implicated may or may not be that he would use a sentence whose role in the language mirrored in some ways the sentence one uses in the language of the report. Plausibly this is so when we report that the ancients discovered that Phosphors was identical to Hesperus. While literally false, its use implicates something true, namely, that the ancients had two words for Venus, which they discovered to codenote. However, if we speak about the beliefs of people who had only one name for Venus, it does not matter which of our names for Venus we use in reporting their beliefs about Venus, and if these facts are known to our audience, we are not taken to implicate anything about the metalinguistic beliefs of the people whose beliefs we report.

23. If this is correct, then we do not need to treat attitude sentences as involving semantically some implicit indexical reference to a mode of presentation or the like ((Schiffer 1977), (Crimmins 1992), (Richard 1990)) or as involving quantification over modes of presentation (Salmon 1986a). Furthermore, this shows that Schiffer's (1987) argument that natural languages are not compositional fails, since his argument depends crucially on the claim that substitutions of coreferring singular terms in attitude contexts do not always preserve truth value.