

Lecture 3

I argued in the previous lecture for a relationist solution to Frege's puzzle, one which posits a semantic difference between the pairs of names 'Cicero', 'Cicero' and 'Cicero', 'Tully' even though there is no semantic difference in the names themselves. My argument for this conclusion was somewhat indirect. I attempted to show that the solution to a variant of Frege's puzzle called for a notion of semantic fact (or requirement) that was not closed under classical consequence and that the semantic difference between the pairs might then be explained in terms of this notion. For the identical pair of names will be *strictly* coreferential, it will be a semantic requirement that the names be coreferential, while the non-identical pair of names will only be accidentally coreferential.

But one might well have thought that if there is a semantic difference between the pairs, it should be evident what it is and should not call for an elaborate argument. I think this is a fair criticism and so let me now attempt to state more directly what that difference is. The names 'Cicero' and 'Cicero' in the identity 'Cicero = Cicero' both (semantically) represent the same object; and so do the names 'Cicero' and 'Tully' in 'Cicero = Tully'. But the first pair of names represents the objects *as the same* whereas the second pair does not.

I take it that we all have an intuitive grip on this notion of representing as the same. But a good test of when an object is represented as the same is in terms of whether one might sensibly raise the question of whether it *is* the same. An object is represented as the same in a piece of discourse if no one who understands the discourse can sensibly raise the question of whether it is the same. Suppose that you say 'Cicero is an orator' and later say 'Cicero was honest', intending to make the very same use of the name 'Cicero'. Then anyone who raises the question of whether the reference was the same thereby betrays his lack of understanding of what you meant.

Other philosophers, besides relationists, can recognize the phenomenon. The relationist understanding of the phenomenon requires two further theses. First, the phenomenon is indeed semantic. When a piece of discourse represents an object as the same, then this is a semantic feature of the expressions by which reference to the object is made. Second, the phenomenon is essentially relational; there are no intrinsic semantic features of the individual expressions in

virtue of which they represent the object as the same.

Both of these further theses might be questioned; and there are two principal lines of attack that we have considered (though there are others). The advocates of 'logical form' argue that the phenomenon is pre-semantic. The difference between the cases is one of logical form; and semantics or representational role does not come into it. But why do we say that the logical form of 'Cicero = Cicero' is 'a = a' rather than 'a = b'? It cannot be because we have the same name on the left and right of the identity; for the name on the left could have been used for the orator and the name on the right for the spy. But then what else could it be other than that the names are used to represent their object as the same? Thus the identity or difference of symbols in the logical form is merely a reflection of the phenomenon of interest to us and cannot itself be taken to be that in which the phenomenon consists.

The more serious challenge is from the intrinsicalist. He agrees that the phenomenon is semantic but argues that when two expressions represent an object as the same it is because they represent that object in the same *way*. Now the relationist can agree that when two expressions represent an object as the same they represent the object in the same way. For, trivially, the first expression will represent the object as the same as the first expression while, by assumption, the second expression will not represent the object as the same as the first expression. However, he will deny that one can *account* for the expressions representing the object as the same in terms of *how* each represents its object. In other words, there is not some intrinsic semantic feature of the expression, the way it represents its object, whose common possession accounts for the two expressions representing the object as the same.

There are two main objections to the intrinsicalist. The first is that his belief in these ways of representing the object is a myth. We might think of them as the ghost of images under the old imagist theory of meaning; these images may have been drained of their psychological content but they still operate as intrinsic aspects of meaning. The other objection is that these ways of representing an object are not, in any case, able to do the job required of them since, even if the ways were the same, there would still arise the question of whether *they* were represented as the same. I shall pursue these objections later, both towards the end of the present lecture and in the final lecture.

This way of defending relationism rests upon (a) making out that there is phenomenon of coordination or representing-as-the-same, (b) showing, contra the advocates of logical form, that is not pre-semantic and (c) showing, contra the advocates of 'ways of representing', that it is not intrinsicalist. It need not rest, as I previously made it rest, upon accepting a certain notion of semantic fact. Thus someone who was dubious about that notion might still be willing to accept the above arguments for adopting a relationist conception of coordination.

I therefore recommend that we view my discussion of the variant puzzle in a somewhat different light. Many referentialist have supposed that they must give up Transparency - the view, roughly speaking, that meaning is accessible to the competent speaker. But what my discussion shows is that the referentialist is able to accept a form of Transparency as long as he willing to reject a form of Closure (the principle that semantic facts are closed under consequence). This is an interest of the discussion that is independent of the issue of relationism; and it is perhaps worth saying a few more words about it.

A standard formulation of Transparency is as follows:

if the meaning of two expressions is the same then this is something that someone who understands the expressions is in a position to know (or, in my terminology, it is 'accessible to the understanding').

But it seems to me that Transparency, as so formulated, is better regarded as the product of two further theses. One which, in this context might be called 'Transparency Proper', is Transparency as I originally formulated it:

Semantic facts are accessible to the understanding.

The other is a thesis about the status of identity of meaning:

That two expressions mean the same is a semantic fact.

I have suggested that the referentialist might hold on to Transparency Proper and yet reject the subsidiary thesis about status. Of course, it is odd to deny that synonymy is a semantic fact. But we should note that there is a stricter notion of synonymy - what we previously called strict coreference - which can plausibly taken to constitute a semantic fact. We should also note that everything we have so far said is compatible with accepting that it is a semantic fact that any given expression means what it does, which, in conjunction with Transparency Proper, will yield

the conclusion that meaning (as opposed to identity of meaning) is also accessible to the understanding.

But once Transparency is defended in this way, it can also be shown to provide us with an independent way of defending relationism. For we may define two expressions to be strictly coreferential when it is a semantic requirement that their reference be the same. For the referentialist, strict coreference will then be an essentially relational notion; whether it holds between two names will not simply be a function of their intrinsic semantic features. It will hold between 'Cicero' and 'Cicero', for example, but not between 'Cicero' and 'Tully', even though their referents, and hence their intrinsic semantic features, are the same.

This independent defense of relationism does not require that we equate the intuitive notion of two names representing an object as the same with the semi-technical notion of strict coreference. However, that the two notions coincide is a plausible view and one that I am inclined to adopt.

The purpose of the present lecture is twofold: to develop a relational semantics for the use of proper names; and to defend relationism for thought. It will be seen that a language containing proper names will submit to a relational semantics that is, in some ways, analogous to the relational semantics for variables; and it will be shown that, just as there are irreducible semantical relationships within language, so there are irreducible representational relationships within thought.


There is a familiar referentialist semantics for a language containing names. Under this semantics, the semantic value, or content, of a name is the individual that it names, the content of a predicate is a property, and the content of a logical connective is an operation on propositions. The content of complex expressions will then be determined in familiar compositional fashion from the contents of the simpler expressions from which it formed. Thus, given that the content of 'Cicero' is a certain individual and the content of 'is an orator' a certain property, the content of the sentence 'Cicero is an orator' will be a proposition to the effect that the individual has the property.

The most distinctive feature of the semantics is that the propositions assigned to sentences


containing names will be *singular*; where the names refer to certain individuals, the propositions will involve those individuals themselves without regard for how they might be identified. In stating the semantics, there are somewhat different conceptions of a singular proposition with which one might work. But all that is important for our purposes is that the individuals should actually occur as constituents in the singular propositions to which they belong. Thus it should make sense to talk of a singular proposition containing one or more occurrences of a given individual; and it should also make sense to talk of replacing certain occurrences of an individual in a singular proposition with others. So, for example, we should be able to say that the singular proposition that Cicero is an orator contains one occurrence of Cicero and that the singular proposition that Cicero is identical to Cicero contains two occurrences; and we should also be able to say that the singular proposition that Bush is an orator is the result of substituting Bush for Cicero in the proposition that Cicero is an orator.

The familiar referentialist semantics ignores semantic relations. Each occurrence of a name within an expression puts an individual into the content that it signifies; and it makes no difference to the content signified whether two names within the expression are strictly or accidentally coreferential. We should now like to develop a semantics that is able to take account of semantic relations; differences in semantic relation between names should be reflected as differences in content.

The natural suggestion is that these differences in 'coordination' among names should be reflected as differences in coordination among the individuals to which they correspond. Consider the sentences 'Cicero admires Cicero' and 'Cicero admires Tully', for example. They express the same singular proposition, one to the effect that a given individual, Cicero, admires Cicero. The proposition in question will therefore contain two occurrences of the individual Cicero. But, in the proposition expressed by the first sentence, the two occurrences of Cicero would be taken to be coordinated, thereby indicating that they were represented as the same, whereas in the proposition expressed by the second sentence, they would be taken to be uncoordinated, thereby indicating that they were not represented as the same. Thus we might depict the first proposition by:



with a connecting line to represent the coordination between the two occurrences of Cicero, and we might depict the second proposition by:



with an annulled line to represent the lack of coordination between the two occurrences of Cicero. Just as there may be coordination (or lack of it) within a proposition, so there may be coordination (or its lack) across propositions or other kinds of content. Consider the sentence-pairs ‘Cicero is Roman’, ‘Cicero is an orator’ and ‘Cicero is Roman’, ‘Tully is an orator’. Each pair of sentences expresses the same pair of singular propositions. But there is a semantic difference between the two pairs, since the subject-terms are strictly coreferential in the first pair, though not in the second. If we are to reflect this difference at the level of the propositions, then we must require that the (identical) subjects of the propositions be coordinated in the first case but not in the second.

The general notion of propositional coordination might be defined as follows. Suppose we are given a sequence of propositions (or contents) $P = p_1, p_2 \dots, p_n$. By a *coordination-scheme* \mathcal{C} for P is meant an equivalence relation on the occurrence of individuals in $p_1, p_2 \dots, p_n$ which is such that two occurrences of individuals are related by \mathcal{C} only if they are occurrences of the same individual. Intuitively, a coordination-scheme tells us how the occurrences of individuals within the propositions are or are not coordinated. A *coordinated sequence of propositions* (or contents) is then an ordered pair $P; \mathcal{C}$, where P is a sequence of propositions (or contents) and \mathcal{C} is a coordination-scheme for P . Recall the notion of a coordinated sequence of open expressions from §2.5. The present notion is an abstract intensional counterpart of the earlier syntactic notion.

A compositional semantics for assigning coordinated content may then be given in a way that it analogous to our previous semantics for variables. To the pair of names ‘Cicero’, ‘Cicero’,

for example, will be assigned the coordinated sequence of individuals Cicero, Cicero while, to the pair of names ‘Cicero’, ‘Tully’, will be assigned the uncoordinated sequence of individuals Cicero, Cicero. The coordinated propositions assigned to ‘Cicero = Cicero’ and ‘Cicero = Tully’ can then be determined in the obvious manner on the basis of the semantic value of ‘=’ and the coordinated sequence of individuals assigned either to ‘Cicero’, ‘Cicero’ or to ‘Cicero’, ‘Tully’. Similarly, to the pair of sentences ‘Cicero is a Roman’, ‘Cicero is an orator’, will be assigned the coordinated sequence of propositions *Cicero is a Roman*, *Cicero is an orator* while, to the pair of sentences ‘Cicero is a Roman’ and ‘Tully is an orator’, will be assigned the uncoordinated sequence of the same propositions. The coordinated propositions assigned to the conjunctions ‘Cicero is a Roman and Cicero is an orator’ and ‘Cicero is a Roman and Tully is an orator’ may then be determined in the obvious manner on the basis of the coordinated sequences of propositions assigned to the respective sequences of their conjuncts. It should be noted that it is not possible to determine the coordinated content of ‘Cicero = Cicero’ on the basis of the separate contents of ‘Cicero’ and ‘Cicero’, since we could not then distinguish between the coordinated contents of ‘Cicero = Cicero’ and ‘Cicero = Tully’; and, similarly, it is not possible to determine the coordinated content of the conjunction ‘Cicero is a Roman and Cicero is an orator’ on the basis of the coordinated content of the individual conjuncts. The semantics is *essentially* relational.

[With the framework of coordination in place, we are now able to provide a rigorous definition of our earlier notion of manifest consequence. Suppose we have a coordinated sequence of propositions. This then corresponds in a natural way to a non-coordinated sequence of propositions; for occurrences of individuals that are coordinated may be replaced by the same individual and occurrences of individuals that are uncoordinated may be replaced by distinct individuals. In effect, we adopt the perspective of the subject and treat uncoordinated occurrences of the same individual as if they were occurrences of distinct individuals. Given a coordinated sequence of propositions p, q, r, \dots , we may then say that p is a manifest consequence of q, r, \dots if p' is a classical consequence of q', r', \dots , where p', q', r', \dots is the noncoordinated sequence of propositions corresponding to p, q, r, \dots . In this way, we reduce manifest consequence on coordinated propositions to classical consequence on noncoordinated

propositions. Suppose now that p, q, r, \dots is a noncoordinated sequence of propositions. We may then say that p is a manifest consequence of q, r, \dots if any coordination-scheme on q, r, \dots can be extended to all occurrences of individuals in p in such a way as to make p a manifest consequence of q, r, \dots . In this way, we reduce manifest consequence on noncoordinated propositions to manifest consequence on coordinated propositions. For example, the inference x F's and x G's to $\exists x(x$ F's and x G's) will not be manifestly valid. For suppose that the two occurrences of x in the premisses are uncoordinated. Then the corresponding noncoordinated inference is from x F's and y G's to $\exists x(x$ F's and x G's), x and y distinct, and this inference is *not* classically valid. On the other hand, the inference from x F's and x G's to $(x$ F's & x G's) will be manifestly valid. For consider any coordination-scheme on the premisses, say one in which the two occurrences of x are uncoordinated. Extend this scheme to the conclusion by coordinating the first occurrence of x in the premisses with the first occurrence of x in the conclusion and the second occurrence of x in the premisses with the second occurrence of x in the conclusion. The corresponding noncoordinated inference (say from x F's and y G's to $(x$ F's & y G's) will then be classically valid.]

The familiar referentialist semantics is one-tier - there is merely one level of semantic value; and this is in contrast to the standard Fregean semantics, which is two tier - it computes semantic value both at the level of reference and at the level of sense. In this respect, the relational version of the referentialist semantics is closer to the Fregean semantics, since it computes semantic value both at the level of noncoordinated and at the level of coordinated content. There are some significant differences between the two approaches, which are worth spelling out.

First, the semantics of sense runs strictly parallel to the semantics of reference, but the semantics of coordinated content does not run strictly parallel to the semantics of noncoordinated content, for the former is relational, it works off semantic connections, while the latter is intrinsicalist, it works off semantic values. Second, under the Fregean account, sense and reference stand in an external relationship; the sense is genuinely determinative of the reference and it is, in general, a contingent matter that the sense has a given referent. Under our referentialist view, on the other hand, coordinated and non-coordinated content stand in an

internal relationship; the non-coordinated content is 'built into' the coordinated content and it is impossible for a give coordinated content to be associated with a different non-coordinated content. Third, sense is not a relational matter. It is to be assigned to single expressions, considered one at a time, and the only meaning we can give to the sense of several expressions is in terms of the sense that is given to each individual expression. Coordination, on the other hand, *is* a relational matter. Coordinated content may be assigned to several expressions, considered together, and it is not, in general, a function of the coordinated content of each expression. Fourth, the sense of a sentence is truth-conditional in character; it bears upon the conditions under which the sentence is true. But the coordinative aspect of the coordinated content of a sentence is purely conceptual in character. It bears, not upon the conditions under which the sentence is true, but upon the conditions under which its primary content might be grasped. Thus the truth-conditions of 'Cicero admires Cicero' and 'Cicero admires Tully' are the same and the difference in their coordinated content relates solely to how those truth-conditions are to be grasped. Finally and most significantly, sense is, in many respects, more fine-grained than coordinated content. The coordinated contents of 'Cicero is an orator' and of 'Tully is an orator' are the same but their sense, for the Fregean, will be different. Thus coordinated content is a poor substitute for sense and, if the work done by sense is to be done by coordination, then it is not simply by using coordinated content in place of sense.

I have so far confined my attention to relationism at the level of language. I now wish to argue that it also holds at the level of thought - that just as there are semantical relationships within language that are not to be understood in terms of the intrinsic semantic features of language, so there are representational relationships within thought that are not to be understood in terms of the intrinsic representational features of thought.

I previously maintained that there was an intuitive notion of two names semantically representing a given individuals as the same. It seems to me that, similarly, there is an intuitive notion of a thought, or thoughts, mentally representing a given individual as the same. Suppose I try to recall what I know about Cicero. You think: *he is a Roman*; and you also think: *he is an orator*. Your two thoughts then represent that individual, in coordinated fashion, as the same.

To represent the individual as the same is not, of course, simply a matter of thinking of the same object. If you would have expressed your first thought in the words ‘Cicero was a Roman’ and your second thought in the words ‘Tully was an orator’, then your thoughts would not have represented that individual *as* the same even though they *are* the same.

I wish to maintain that, just as the semantical form of coordination is essentially relational, then so is the mental form. But since new arguments apply in this case, it will be worth considering the matter anew.

Many philosophers, I suspect, will be tempted to explain mental coordination in terms of ‘mental files’. It is supposed that we keep mental files of every individual of which we are capable of having a singular thought. To represent an individual as the same is then to associate the same mental file with the individual. Thus in the example above, the individual Cicero, in the two thoughts, will be associated with the same mental file.

In meeting this objection to relationism, we may grant that representing an individual as the same in thought consists in associating the same mental file with the individual; and we may also grant that this is an intrinsic feature of the thought. However, representational relationism is not the view that representational relationships are not intrinsically grounded. It is the view that representational relationships are not intrinsically grounded in *representational* features. Our interest is solely in the representational realm; and our question is whether, within this realm, there are aspects of thought that are essentially relational.

Now it seems clear that the mental file associated with a given individual in a thought is not a representational aspect of the thought. In having the thought, we are not thinking of the mental file and nor are we thinking of the individual as an individual with that mental file. The mental file operates below the representational surface, so to speak, and is not itself part of the content of what is thought.

There are two other alternatives to relationism that the referentialist might suggest. According to the first, when one represents an individual as the same in thought, one is having the additional thought that the individuals are the same. Thus what it is to think that the individual Cicero is a Roman and then to have the coordinated thought that he is an orator is to

think that the individuals are the same (or to intend that they be the same). But if the additional thought is to have the desired effect, we must already presuppose that the individuals in that thought are represented as the same as the individuals in the given thoughts. Unless there is a coordination between the additional thought and the original thoughts, there will not be coordination between the original thoughts. Thus this view already presupposes what is in question. Nor will it do to suggest that the additional thought is that the individual of the first thought is identical with the individual of the second thought, since that would make the thoughts themselves part of the content of what is thought.

One should be careful not to conflate representing two objects as the same with representing them as *being* the same, i.e. with having the thought that they are the same. Indeed, far from being identical, these two relationships are largely incompatible, for I can only significantly represent two individuals as *being* the same if I do not already represent them *as* the same. A common manifestation of this conflation is in the discussion of mental files. It is supposed that when I learn, for example, that Cicero is Tully, the two mental files associated with 'Cicero' and 'Tully' are merged into a single mental file. There is a double incoherence here. First, the view must be that, in representing the individual as being the same, I thereby represent them as the same since, otherwise, there is no reason to think that a merger would automatically follow. But as I have already pointed out, representing the individual as being the same, in this case, requires that it not be represented as the same. Second, talk of merger is out of place. It is not that the merged file represents the individual as the same as the earlier files, since that would require that the given files represent the individual as the same. Rather, the new file, if one chooses to create it, will represent the individual as being the same as the previous files. One *copies* the contents of the previous files to the new file; one does not simply *combine* them.

The other referentialist alternative to relationism is not to add a thought to the original content but to revise one's view of what that content is. Thus to have the thought that one might express in the words 'Cicero = Tully' is to have the thought that Cicero is identical to Cicero while to have the thought that one might express in the words 'Cicero = Cicero' is to have the thought that Cicero is self-identical.

There are several objections to this view. But the main one is that it is not able to deal

with coordination across different kinds of thought. Suppose I judge that Cicero is a Roman and then I make the coordinated judgement that Cicero is an orator. Then, at a stretch, I might regard this as a compound judgement to the effect that Cicero is a Roman orator. But suppose that I *wonder* whether Cicero is an orator and then make the coordinated judgement that he is an orator. The current view then requires us to take this as a single 'thought' directed at a single content. But what is the thought and what is the content? There is nothing sensible that we can say.

Finally, we should reconsider the standard intrinsicist response to the relationist view. When an individual is represented as the same in thought, according to this view, it is because the individual is represented in the same way; coordination is achieved through coincidence in 'mode of presentation'. The difficulty with this view is to say in particular cases what the relevant modes of presentation might be. Imagine a variant of the 'Paddy' case. I see someone and call him 'Paddy'. I see someone on another occasion and also call him 'Paddy'. I think the two people are different when, in fact, they are the same. I therefore have two uses of the name 'Paddy', which we might dub 'Paddy1' and 'Paddy2'. Thoughts mediated through the use of 'Paddy1' or through the use of 'Paddy2' will be coordinated with one another, while thoughts mediated through the use of 'Paddy1' and 'Paddy2' will not be; and so the view requires that the various uses of 'Paddy1' or of 'Paddy2' should be associated with the same mode of presentation, while the different uses of 'Paddy1' and 'Paddy2' should be associated with different modes of presentation.

But what is the difference in mode of presentation? It cannot lie in the normal descriptive material that is associated with the two uses of the name, since that may be the same. Indeed, I might take the individuals to be identical twins and attribute the very same properties to them both. In the absence of a descriptive difference, some philosophers have been tempted to appeal either to a difference in the use of the names themselves or to a difference in their origin, since one use of the name is connected with one encounter with Paddy and the other use with another encounter. But both responses are objectionable in that they fail to point to a representational aspect of thought. When I think about the individual through the use of 'Paddy1' or 'Paddy2', I am not thinking about the use of the name or my encounter with the individual and nor am I thinking of the individual as one that was named or encountered in a certain way. Indeed, I might

have completely forgotten about my original encounters with Paddy and yet still be able to make continued use of 'Paddy1' and 'Paddy2'. The account in terms of names is also circular. For part of what constitutes the use of a name is that the different manifestations of that use should be coordinated. Thus one cannot even say what the two uses of 'Paddy' are without appeal to the notion of coordination.

It therefore appears that coordination can sometimes be achieved without coincidence in mode of presentation. But if it can *sometimes* be so achieved, then this suggests that it is *always* so achieved. For what could stand in its way? Thus the presence or absence of a common mode of presentation emerges as entirely incidental to whether or not there is coordination.

Given a relationist view of thought, we should now take the content of a thought to be a coordinated proposition. Thus when I think to myself, *Cicero loves Cicero*, the content of my thought is the coordinated identity proposition while, when I think to myself, *Cicero loves Tully*, the content of my thought is the uncoordinated identity proposition. Similarly, the content of a *set* of thoughts should be taken to be a coordinated set - or system - of propositions. (Up to now, we have only defined the notion of a coordinated *sequence* of propositions. But a coordinated *system* might be taken to be the result of 'factoring out' the order from a coordinated sequence.)

There are two consequences of this relational understanding of content that will later be of great interest to us. Suppose we ask what someone believes. Then under the relational approach it is not sufficient simply to specify the content of each of his beliefs. Suppose that A believes that Cicero is a Roman and also has the coordinated belief that Cicero is an orator, while B believes that Cicero is a Roman and also has the uncoordinated belief that Cicero is an orator. Then the contents of their individual beliefs are the same; and yet the content of their beliefs is different, since the two propositions are coordinated for A though not for B. Thus what someone believes is not simply a function of his individual beliefs.

It also follows on the relational approach that we cannot adopt a simple incremental view of what it is to acquire knowledge or new beliefs. Suppose that A and B both believe that Cicero is a Roman. Each now learns that Cicero is an orator. But A learns the proposition in a way in which it is coordinated with his previous beliefs while B learns it in a way in which it is uncoordinated with his previous beliefs. Then even though they both begin by believing the

same propositions, and even though they both learn the same proposition, they end up believing something different, since A now has a belief in a coordinated system while B has a belief in an uncoordinated system of propositions.

In the next lecture, I wish to discuss the connection between language and thought; and I shall be especially interested in the question of how we might communicate our thoughts by means of language. I shall argue that, not only is there coordination within language or within thought, but that there is also coordination between language and thought and that it is essential for understanding the nature and possibility of communication that we adopt the relational point of view.