

Lecture 2

I argued in the previous lecture that a satisfactory semantics for variables should be relational; it should presuppose that there are semantic relationships between variables that are not grounded in their intrinsic semantic features. The critical point in favor of such a view is that a complete description of the semantic behavior of variables is not given by stating the range of values that each variable can assume. It must also include a statement of how the value assumed by one variable may or may not be capable of constraining the values assumed by other variables. This is essentially a relational matter and cannot be read off from the semantic behavior of each individual variable.

The topic of the present lecture is Frege's puzzle. I shall argue that just as the antinomy of the variables requires us to adopt a relational account of variables, so Frege's puzzle requires us to adopt a relational account of names. However, before embarking on this topic, we should complete a task left over from the previous lecture. Given a relational view of variables, then how should the semantics for first-order logic proceed?

The aim of the standard, intrinsicalist semantics is to assign a semantic value to each (meaningful) expression of the language under consideration. Suppose that an expression E is syntactically derived from the simpler expressions E_1, E_2, \dots, E_m . Then the semantic value $|E|$ of E is taken to be the appropriate function $f(|E_1|, |E_2|, \dots, |E_m|)$ of those simpler expressions. Given semantic values for the lexical items of the language (those not derived from other expressions), the semantic value of each expression is then determined.

The aim of a relational semantics, by contrast, is to assign a semantic connection to each sequence of expressions. Such a connection is intended to encapsulate not only the semantic features of each individual expression but also the semantic relationships between them. The semantic value $|E|$ of an expression E is then taken to be a function $f(|E_1, E_2, \dots, E_m|)$ of the semantic connection $|E_1, E_2, \dots, E_m|$ of the expressions from which it is derived and, in general, the semantic connection $|F_1, F_2, \dots, E, \dots, F_{n-1}, F_n|$ on the sequence $F_1, F_2, \dots, E, \dots, F_{n-1}, F_n$ is taken to be a function $f(|F_1, F_2, \dots, E_1, E_2, \dots, E_m, \dots, F_{n-1}, F_n|)$ of the semantic connection on the simpler sequence $F_1, F_2, \dots, E_1, E_2, \dots, E_m, \dots, F_{n-1}, F_n$. Given semantic connections on sequences of lexical semantics, the semantic connection on any sequence of expressions is then determined.

Such a semantics is still compositional. Indeed, it is still correct to say that the meaning of an expression is to be determined on the basis of the meaning of its components. But the phrase the 'meaning of its components' must now be taken in a collective sense to include the meaning relationships among the components and not just their individual meanings.

If we are to apply the general idea of a relational semantics to the language of first-order logic, we must specify the syntax upon which the semantics will be based. For the most part, this is standard: a disjunctive formula $(A \vee B)$, for example, will be derived from the disjunction operator ' \vee ' and the disjuncts A and B , and an atomic sentence $Pt_1t_2 \dots t_n$ will be derived from the predicate P and the argument-terms t_1, t_2, \dots, t_n . However, in two key respects our treatment of the syntax is not altogether standard. In the first place, the lexicon will now be taken to include variables, to which semantic values or connections should be assigned. In the second place, the quantified expression $\exists xA$ will be taken to be derived from the quantifier \exists , the bound variable x , and the embedded formula A . Thus the bound variable x comes into its own as one of the syntactic constituents of the whole formula.

So much for the syntax. In order to set up the semantics, we need some conception of semantic connection. This may be obtained by generalizing the notion of a value-range for a variable. The value-range of a variable is the set of values it is capable of assuming; and similarly, given any sequence of expressions, we may take its value-range - or semantic connection - to be the set of sequences of values that those expressions are simultaneously capable of assuming.

The lexical semantics is, for the most part, straightforward: extensions should be assigned to predicates, denotations to constants, and functions to function symbols. However, we now include variables within the lexicon and so the lexical semantics should also specify the semantic connection on any sequence of variables. Suppose that we are given the sequence of variables x, y, x, y , for example. Then in conformity with our understanding that distinct variables take values independently of one another and that identical variables take the same value, the semantic connection on this sequence should be the set of all quadruples a, b, c, d of individuals from the domain for which $a = c$ and $b = d$; and similarly in the more general case. We might call this the 'coordination rule'.

We also need rules for extending the semantic connections to more complicated expressions and more complicated sequences of expressions. Consider, by way of example, the complex terms $x.x$ and $x.y$. The first should have as its value-range the set of all non-negative reals; and the second should have as its value-range the set of all reals whatever. How do we secure this result? If we let the value-range of $x.x$ simply be a function of the value-range of x and x , and similarly for $x.y$, then we cannot distinguish between them, since the value-ranges of x and y are the same. However, we take the value-range of $x.x$ to be a function on the semantic connection on x, x

and the value-range of $x.y$ to be a function on the semantic connection on x, y . These semantic connections differ, as we have seen; the first comprises all identical pairs of reals, while the second comprises all pairs of reals whatever. And there is therefore a corresponding difference in the value-ranges of $x.x$ and $x.y$; for each will comprise the corresponding set of products and will thereby yield the desired result.

Quantifiers introduce complications. Suppose we wish to evaluate the sequence of formulas $\exists x(x > 0), x$. Then we will want to explain the semantic connection on $\exists x(x > 0), x$ in terms of the corresponding semantic connection on $x, x > 0, x$ (or, more directly, on $\exists, x, x > 0, x$). But now the two occurrences of x that derive from the quantified formula will become accidentally bound to the other occurrence of x .

How are we to avoid this result? If we were to evaluate the sequence $x, x > 0, x$ without regard for its syntactic origin in the sequence $\exists x(x > 0), x$, we would take all three occurrences of the variable x to be bound. We must therefore somehow distinguish between the two cases.

What I would like to propose is that we explicitly indicate whether or not two variables are to be 'coordinated'. Thus in the case in which $x, x > 0, x$ is derived from $\exists x(x > 0), x$, we will take the first occurrences of the variable to be 'coordinated' with one another though not with the third; and in the case in which it is not so derived, all three occurrences of the variable will be coordinated. We can now evaluate quantified formulas without falling prey to accidental binding.

Modest as this proposal might appear to be, its development calls for a fundamental revisions in our previous formulation of the semantics. In the first place, the syntactic object of evaluation will no longer be a sequence of expressions but a *coordinated* sequence of expressions. This is a sequence of expressions E_1, \dots, E_n along with a *coordination-scheme* \mathcal{C} which tells us when two occurrences of the same variable are to be coordinated (formally, a coordination-scheme is an equivalence relation on the occurrences of variables in the sequence which only relates occurrences of the same variable). In the second place, the coordination rule must be modified. Instead of requiring that all occurrences of the same variable should receive the same value, we should only require that they receive the same value when they are coordinated.

There is a way of making the present proposal more familiar and perhaps more palatable. Several philosophers have remarked that a 'telegraphic' notation - one in which lines replace variables - would somehow be more perspicuous than the standard notation. But, unfortunately, they never attempted to make the notation precise or to say how a semantics directly imposed upon such a notation might go. But once this is done, it will be seen to lead to something like the present proposal. Thus in so far as the instincts of these philosophers were on the right track, they provide further support for our view.

We have seen that the doctrine of semantical intrinsicism fails for variables. I wish to argue in the remaining lectures that the doctrine also fails for other, more familiar, parts of speech; and, as with the case of variables, we shall attempt to establish this conclusion through consideration of various philosophical puzzles. In each case, it will be argued that relationism provides the most plausible way in which the puzzle might be solved.

_____ We begin with Frege's famous puzzle (though I make no claims of historical accuracy). Take two coreferential names, say 'Cicero' and 'Tully'; and consider the identity-sentences, 'Cicero = Cicero' and 'Cicero = Tully'. Say that two sentences are *cognitively different* if they can convey different information to someone who understands both sentences; say that two meaningful expressions are *semantically different* if they differ in their meaning or 'semantic role'; and say that two referring expressions are *referentially different* if they are not coreferential. The puzzle might then be seen to be based upon the following five assumptions:

- (1a) Cognitive Difference The two identity sentences are cognitively different;
- (1b) Cognitive Link If the sentences are cognitively different, then they are semantically different;
- (2) Compositionality If the sentences are semantically different, then the names 'Cicero' and 'Tully' are semantically different;
- (3) Referential Link If the names 'Cicero' and 'Tully' are semantically different, they are referentially different;
- (4) Referential Identity The names 'Cicero' and 'Tully' are not referentially different.

The five assumptions are jointly inconsistent; and so at least one of them should be given up. The puzzle is to say which and why.

There is an abridged version of the puzzle in which the Cognitive Difference and Link is replaced by the following consequence of them:

- (1) Semantic Difference The two identity sentences are semantically different;
- We begin with the abridged, purely semantical version of the puzzle and only later consider the complications that arise from the cognitive aspects of the puzzle.

There have been two main responses to the puzzle - the Fregean and the referentialist. Both sides accept

Compositionality and Referential Identity. The only alternatives are therefore to reject Semantic Difference (there is a semantic difference in the identity sentences) or Referential Link (no semantic difference without a referential difference). The Fregeans reject the Referential Link but accept Semantic Difference: they maintain that, even though the reference of the names 'Cicero' and 'Tully' is the same, their meaning or 'sense' is different; and given that *their* meaning is different then so, plausibly, is the meaning of the identity-sentences. The referentialists, on the one hand, reject Semantic Difference but accept the Referential Link: for them, there is no more to the meaning or semantic role of a name than its referent; and given that the meaning of the names is the same then so, plausibly, is the meaning of the identity sentences.

Let me briefly review some of the considerations for and against these two responses. The endorsement of Semantic Difference is a strong point in favor of the Fregean response. We have an intuitive notion of meaning and it seems evident that, under this intuitive notion, the two identity-sentences differ in their meaning. Indeed, the difference is not even of a slight or subtle sort; and it is a major mark against the referentialist view that it cannot respect these strong and striking intuitions.

There is also a strong *argument* in favor of Semantic Difference. For even if the intuitive evidence in its favor is rejected, it barely seems possible to reject the intuitive evidence in favor of Cognitive Difference; for surely one may learn something different upon being told 'Cicero = Tully' and upon being told 'Cicero = Cicero'. But it is hard to see how to account for this possible cognitive difference except in terms of a semantic difference.

The main problem with the Fregean position, to my mind, is to say, in particular cases, what the difference in the meaning or sense of the names might plausibly be taken to be. For it seems perfectly possible to have the use of two coreferential names - such as 'Cicero' and 'Tully' - and yet attach the same sense to them both.

The Fregeans have been very resourceful in coming up with possible differences in sense for the problem cases that have been raised against them. They have appealed, for example, to the sense that others might attach to the name or to something meta-linguistic, like *the referent of this name*. And so it may be worth mentioning a case that would appear to be resistant to counter-moves of this sort. The inspiration for the case is our previous example of a mini-semantic-universe in which there are two variables that are intrinsically the same and yet relationally different. Given the existence of such an example for variables, one naturally wonders whether there might not exist a similar kind of example for names.

To this end, let us imagine a universe that is completely symmetric around someone's center of vision. Whatever she sees to her left *is* and *looks* qualitatively identical to something she sees on her right (not that she conceptualizes the two sides as 'left' and 'right' since that would introduce an asymmetry). She is now introduced to two identical twins, on her left and on her right, and she simultaneously names each of them 'Paddy'; using a left token of 'Paddy' for the left twin and a right token of 'Paddy' for the right twin. The two tokens of 'Paddy' are then used in tandem so as not to disturb the symmetry. Thus if she uses a left token of 'Paddy' to say 'Paddy was wearing pink pants', she simultaneously uses a right token of 'Paddy' to say the same thing. She can even assert the non-identity of the two Paddies by simultaneously uttering the one token of 'Paddy' from the left side of her mouth, the other token from the right side, and a word for non-identity from the very middle of her mouth.

It seems clear that she has the use of two names (and if one wanted their referents to be the same, we might suppose that the uses originated, not in two identical twins, but in two identical images of the same person). But what is the difference in sense? By symmetry, there is no purely descriptive difference in the referents. And this in itself is enough to refute a view that takes sense to be a purely descriptive means of identifying a referent.

But what of a more liberal view of sense? Given that our subject 'picks out' the two individuals in two different ways, then might this not be taken to constitute a difference in sense? But what exactly are these different ways of picking out the individuals meant to be? There would appear to be only two plausible candidates. They could be ways in which the individuals are *currently* picked out; the sense of a token of a name, in other words, would somehow be tied to the use of that very token. But in this case, the sense of the name would vary from one moment to the next; and yet surely this is not so - or, at the very least, surely it should be possible for our subject to use two consecutive tokens of the name with the very same sense. The other alternative is to look at the ways in which the two names were *originally* picked out; the sense of each token of the names would then be tied to the original identification of the individuals. The problem here is that it would appear to be compatible with the continued use of each name that the subject should irretrievably lose all knowledge of how its referent was originally identified; and, in this case, she would be put in the odd situation of being able to use a name with a sense of which she was incapable of having any knowledge. Thus in order to evade the difficulty, the Fregean must give up either the reproducibility or the accessibility of meaning; and neither option seems particularly attractive.

In current philosophical thinking we seemed to have reached an impasse on the puzzle, with strong arguments in favor of both Semantic Difference and Referential Link and yet no way to choose between them. And this suggests that we should perhaps take more seriously the possibility of rejecting the assumption of Compositionality

that puts them in conflict. For we might then affirm both that there is no semantic difference between coreferential names, thereby securing the benefits of the referentialist position, and that there is a semantic (or cognitive) difference between the identity-sentences, thereby securing the benefits of the Fregean position. A more acceptable form of referentialism might thereby be embraced, not subject to the usual Fregean strictures.

But how might Compositionality plausibly be rejected? It is at this point, I believe, that our previous considerations concerning the semantical role of variables may prove helpful. For consider the analogue for variables of the preceding puzzle for names:

(1)' The identity-formulas ' $x = x$ ' and ' $x = y$ ' are semantically different (have a different semantical role);

(2)' If the identity-formulas are semantically different, then so are the variables ' x ' and ' y ';

(3)' The variables ' x ' and ' y ' are not semantically different.

As before, the assumptions are jointly inconsistent; and so one of them must be given up.

However, in the present case, neither the Fregean nor the referentialist response is at all plausible. We can hardly toe the referentialist line by denying (1)', for clearly there is a difference in semantical role between the formulas ' $x = x$ ' and ' $x = y$ '. But nor can we toe the Fregean line by denying (3)', for in what could the semantic difference between the variables ' x ' and ' y ' consist? It is not as if the variables ' x ' and ' y ' have a special ' x '-sense or ' y '-sense not possessed by the other. Thus it looks in this case as if the only reasonable option is to reject the compositionality assumption.

Our relational approach to variables also helps to make plausible *how* Compositionality might be rejected. For we should distinguish between what one might call 'Compositionality Proper' and 'Intrinsicality'.

Compositionality Proper states:

(2)' (a) If the identity formulas ' $x = x$ ' and ' $x = y$ ' are semantically different, then so are the pairs of variables x, x and x, y .

Intrinsicality, on the other hand, states:

(2)' (b) If the pairs x, x and x, y are semantically different, then so are the variables x and y .

Compositionality, as it is usually stated, is the product of Compositionality Proper and Intrinsicality. But we may give up Intrinsicality without giving up Compositionality Proper; and as long as we have Compositionality Proper, we will still have the option of providing a relational semantics for the use of variables. Thus the rejection of Compositionality does not require that we reject the general idea of a compositional semantics; and so the principal reason for adhering to Compositionality is removed.

This suggests that we might be able to say something similar in the case of names. Compositionality Proper now takes the form:

(2)(a) If the identity-sentences ' $\text{Cicero} = \text{Cicero}$ ' and ' $\text{Cicero} = \text{Tully}$ ' are semantically different, then so are the pairs of names ' Cicero ', ' Cicero ' and ' Cicero ', ' Tully ',

while Intrinsicality takes the form:

(2)(b) If the pairs of names ' Cicero ', ' Cicero ' and ' Cicero ', ' Tully ' are semantically different then so are the names ' Cicero ' and ' Tully '.

And so it looks as if we might reject Compositionality by rejecting Intrinsicality without thereby giving up the idea of a compositional semantics for our use of names.

Unfortunately, the analogy with variables will only take us so far. For this strategy requires that, even though there be no semantic difference between the names ' Cicero ' and ' Tully ', there should be a semantic difference between the pairs of names ' Cicero ', ' Cicero ' and ' Cicero ', ' Tully '. There should, in other words, be semantic relationship that holds between ' Cicero ' and ' Cicero ' and yet does not hold between ' Cicero ' and ' Tully '. But what might this relationship be? This is the fundamental difficulty we face in extending the relational approach from variables to names. In the case of variables, we could appeal to the fact that the variables ' x ' and ' x ' take 'coordinated' values whereas the variables ' x ' and ' y ' take their values independently of one another. But in the case of names, the semantic role of each coreferential name is fixed by its referent and so talk of 'coordination' or 'independence' would appear to be out of place.

I now wish to argue that this difficulty can be overcome, but my approach will be somewhat indirect. I shall formulate a new version of Frege's puzzle, one which brings certain cognitive aspects of the puzzle to the fore; and it will be the solution to this puzzle which will then reveal how we might provide a relational solution to the original puzzle. I think of the new puzzle as being addressed to the referentialist and so we shall feel free to make assumptions that only he would accept.

The formulation of the new puzzle will require two terms of art which I should now explain. I want, first, to be able to talk of certain things being *semantic facts*. It is a semantic fact, for example, that the sentence ' Snow is

white' is true or that 'the author of Waverley' designates Scott or that 'Cicero' refers to a particular man, Cicero.

But if this notion of semantic fact is to be properly understood, it is important to distinguish it from some related notions. In the first place, we should distinguish between facts that are semantic as to *topic* and semantic as to *status*. Certain properties and relations are, in a clear sense, semantic; they pertain to the meaning of the expressions to which they apply. So, for example, *truth* is a semantic property of sentences, *designation* a semantic relation between a term and an individual, and *synonymy* a semantic relation between two expressions. A fact may be said to be *semantic* in the topic-oriented sense if it pertains to the exemplification of semantic properties or relations. Thus the fact that 'the author of Waverley' designates Scott or that 'bachelor' is synonymous with 'unmarried man' will be semantic in this sense.

However, within the facts that are semantic as to topic, we may distinguish those that are also semantic as to status. These are the facts that are not merely formulable in semantic terms but also belong to the semantics of a given language (or given languages). Thus the fact that the sentence 'snow is white' is true will not be semantic in this sense since it is not a fact about the semantics of English while, presumably, the synonymy of 'bachelor' and 'unmarried man' will be semantic in this sense.

What does it take for a fact that is semantic as to topic to be semantic as to status? A natural criterion, though not one altogether free of circularity, is that a fact semantic as to status will be wholly consequential upon the meaning of the expressions which it concerns while one not semantic as to status will be partly consequential upon non-semantic considerations. Thus the truth of 'snow is white' might be seen to be consequential upon the following two facts:

- (i) 'snow is white' is true iff snow is white;
- (ii) snow is white.

The first of these is purely semantic but the second is not; and, for this reason, the truth of 'snow is white' will not be semantic as to status. On the other hand, there is no similar 'factoring' of (i) and so its truth *will* be semantic as to status.

In the second place, it will be important to distinguish between semantic *facts* and semantic *truths*. Semantic facts are *propositions* while semantic truths are *sentences*. Thus the proposition that 'Cicero' refers to the particular individual Cicero is a semantic fact, while the sentence "'Cicero' refers to Cicero' is a semantic truth. The semantic fact involves the name 'Cicero', the particular individual Cicero, and the relation of referring, while the semantic truth involves the quotation-mark name "'Cicero'" for 'Cicero', the name 'Cicero' itself, and the predicate 'refers to'. The whole point of what I say will be lost unless one carefully bears in mind that I am talking about the semantic facts themselves and not their linguistic expression.

The formulation of the new puzzle will also call for a certain conception of semantic knowledge. Say that a fact concerning a language (or portion of language) is *accessible to the understanding* if any rational and reflective person who understands the language is thereby in a position to know that the fact obtains. For example: anyone who understands the expressions 'bachelor' and 'unmarried man' is in a position to know that they are synonymous; and anyone who understands the sentence 'snow is white' is in a position to know that 'snow is white' if and only if snow is white. These facts are therefore accessible to the understanding.

Now for the puzzle. It is based upon the following five assumption:

Principles:

Referentialism. It is semantic fact that a proper name refers to the particular object that it does.

Closure. Logical consequences of semantic facts are semantic facts.

Transparency. Semantic facts are accessible to the understanding.

Data:

Cognitive Datum. The fact that 'Cicero' and 'Tully' corefer is not accessible to the understanding.

Semantic Datum. The proper names 'Cicero' and 'Tully' corefer.

These five assumptions are jointly inconsistent. For by Semantic Datum, there is an object *x* to which the names 'Cicero' and 'Tully' both refer. By Referentialism, it is a semantic fact that 'Cicero' refers to *x* and also a semantic fact that 'Tully' refers to *x*. It is a logical consequence of the fact that 'Cicero' refers to *x* and of the fact that 'Tully' refers to *x* that 'Cicero' and 'Tully' corefer. So by Closure, it is a semantic fact that 'Cicero' and 'Tully' corefer. But then by Transparency, it follows that their coreferentiality is accessible to the understanding - contrary to the Cognitive Datum.

Given that the assumptions are inconsistent, one or more must be given up; and, as always, the puzzle is to say which and why. I assume that the two data are beyond reasonable doubt. Certainly, 'Cicero' and 'Tully' (in their familiar use) corefer; and it seems evident that someone who has the use of the two names need be in no position to know that they corefer.

This leaves the three principles. The referentialist should certainly accept Referentialism, since it is a

cardinal tenet of his view that the meaning of a proper name is its bearer. And so the choice is between Transparency and Closure.

Most referentialists - at least in so far as they have considered the matter - have been tempted to reject Transparency. But it seems to me that it is Closure that should be rejected; and so perhaps it is worth making a few remarks as to why I do not consider the rejection of Transparency to be such a plausible option.

There are of course some quite general considerations concerning the nature of language that favor the endorsement of Transparency. I am generally sympathetic to these considerations but I wish, for now, to set them aside and to take up certain considerations that strictly relate to the use of Transparency in the present context. There are three applications of Transparency that are required for the argument of the puzzle to go through. The first of these are to the semantic fact that each of the names 'Cicero' and 'Tully' refers to a particular individual. From this it is inferred by Transparency that these facts are accessible to the understanding. The third is to its being a semantic fact that the names are coreferential. From this it is then inferred by Transparency that their coreferentiality is accessible to the understanding. The first two applications, to the reference of the names, are relatively unproblematic. It is very plausible for the referentialist to maintain that, in having the use of a name, one know to what it refers. But even if he does not think this is generally true, we can so rig the case that the circumstances are as favorable as they might reasonably be for the person who uses the name to have the required *de re* knowledge. The issue is therefore over the application to coreferentiality. The coreferentiality of the two names follows from each name referring to what it does. So what are we to make of this? Should we affirm that it is a semantic fact that they are coreferential and thereby reject Transparency? Or should we deny that it is a semantic fact that they are coreferential and thereby reject Closure?

The issue is complicated by the circumstance that if there is a notion of semantic fact for which Transparency holds there will also be one for which Closure holds. For given such a notion, we may take a fact to be semantic *in the broad sense* if it is a consequence of the semantic facts in the original sense. Semantic facts in the broad sense will then conform to Closure by definition and hence will fail to conform to Transparency. Likewise, if there is a notion of semantic fact for which Closure holds there will be one for which Transparency holds. For given such a notion, we may take a fact to be semantic *in the narrow sense* if it is a semantic fact in the original sense that is accessible to the understanding. Semantic facts in the narrow sense will then conform to Transparency and hence will fail to conform to Closure. This might suggest that the issue is purely verbal.

However, I do not believe that we are entitled to play fast and loose with the notion of semantic fact in this way. We want our notion of semantic fact to be faithful to our intuitions about when there is a semantic difference. But the broad notion is not faithful to these intuitions. This may be illustrated with an apocryphal story about Carl Hempel, the distinguished philosopher of science. When Hempel moved to Princeton, some of the philosophers there found the name 'Carl' too Germanic for their taste, and so they decided to use the English name 'Peter' instead. It is not that they *re-christened* Hempel with the name 'Peter'; rather, they decided to use the name 'Peter' as a *substitute* for the name 'Carl'. Consider now a different scenario in which 'Peter' is introduced, not as a substitute for the name 'Carl', but as a name in its own right: on arriving in Princeton, Hempel is re-christened 'Peter' (perhaps even in ignorance of his German name). For the referentialist, the semantic facts in the broad sense for the two languages or idiolects are the same; 'Peter' refers to Hempel, 'Carl' refers to Hempel, and both corefer. But intuitively, there is a semantic difference between the two languages, since it is a *convention* of the first language, though not of the second, that the name 'Peter' should be coreferential with 'Carl'. Someone who failed to recognize that the two names were coreferential would thereby display his lack of understanding of the first language, though not of the second. But once we work with the narrow notion of semantic fact, the intuitive semantic difference between the two languages can be respected.

This suggests that, under a proper understanding of 'semantic fact', it is Closure rather than Transparency that should be given up. However, if this proposal is to be sustained, two important objections to it must be met. The first is that it is unmotivated. The rejection of Closure may save us from certain awkward consequences; but one would like to have some independent reason for thinking that the notion of semantic fact is one for which Closure should fail. The second objection is that the proposal is unworkable; for we *need* a notion of semantic fact that is subject to Closure. Partly this is because we will later want to say that it is a semantic fact that 'Cicero' is coreferential with 'Cicero'; and, since this semantic fact is consequential upon others, its status as a semantic fact would appear to require some form of closure. But it is also partly because we will want to work with a conception of semantic fact that is capable of doing justice to the compositional character of semantics; and this requires that it be possible to *derive* the semantic facts concerning complex expressions from the semantic facts concerning simpler expressions, thereby preserving their status *as* semantic facts.

I do not believe that the answer to these concerns has anything to do with semantics as such, but must be related to the general question of what the referentialist should say about the possibility of inferential knowledge.

And so let us address this more general question first, before considering how the answer might apply to the case of semantics.

Consider an ideal cognizer, one who is perfectly competent in drawing inferences from what he knows. One then naturally assumes that he will know (or be in a position to know) every classical consequence of what he knows. But this is not something that the referentialist is likely to grant. For our ideal cognizer may know that the object x F's under one 'take' on x and that x G's under another 'take' on x , but not be in a position to infer that x both F's and G's, or even that *something* both F's and G's. He may know, for example, that Paderewski is a brilliant pianist (having heard him at a concert) and also that he is a charismatic politician (having observed him at a political rally), without realizing that it is the same person who is both. Since x 's being both F and G is a classical consequence of x 's being F and x 's being G, the referentialist cannot take knowledge, even for an ideal cognizer, to be closed under classical consequence.

However, the knowledge of the ideal cognizer is far from being logically inert, even when singular propositions are in question. From the fact that x F's, for example, he may infer that something F's; and from the fact that every F G's and x F's, he may infer that x G's. Let us say that a given proposition p is a *manifest* consequence of other propositions q, r, \dots if it is a classical consequence of them and if, in addition, it would be manifest to any ideal cognizer who knew the propositions q, r, \dots that p was indeed a classical consequence of those propositions. In other words, the ideal cognizer would not be handicapped by his different 'takes' on the objects occurring in q, r, \dots in recognizing that p was a classical consequence. Knowledge of the ideal cognizer would then be closed under manifest consequence, even if not under classical consequence.

I shall later provide a more formal account of manifest consequence, one stated entirely in logical terms, but for now let us simply accept that there is such a relation and take our ideal cognizer as providing a test of when it holds.

Corresponding to the distinction between classical and manifest consequence is a distinction between two kinds of *domain*. Just as a theory may be taken to be a class of sentences closed under consequence, so we may take a domain to be a class of *propositions* that is closed under consequence. We then have two kinds of domain, depending upon whether consequence is taken to be classical or manifest. These two kinds of domain correspond to two kinds of attitudes one might have towards its content. On the one hand, one might regard the domain as a possible domain of *facts*. It is then only natural that it should be taken to be closed under classical consequence since the classical consequences of facts are also facts. On the other hand, one might think of the domain as a possible domain of *information*; it consists of what one might know rather than what might be true. It is then only natural that it should be taken to be closed under manifest - rather than classical - consequence, since it is only the manifest consequences of known facts that need be known.

This means that, in the special case of semantics, we may have two different conceptions of the domain of true semantic propositions. We may regard it as a domain of facts that is to be 'found' in the world or as a domain of information that is to be 'found' in the speaker. Our general standpoint is that a semantics for a given language is something that the speaker might know and, once we adopt this standpoint, it then becomes clear that semantics, like any other domain of information, can only be expected to be closed under manifest consequence.

Nor is the resulting conception of semantics unduly restrictive. For, it will be a manifest consequence of the fact that 'Cicero' refers to what it does that 'Cicero' is coreferential with 'Cicero' (this is an inference the ideal cognizer might make). And, in addition, though this requires argument, closure under manifest consequence will permit the kind of derivations required by compositionality. The answer to the two objections is therefore already implicit in any reasonable account of referentialist epistemology.

We are now in a position to provide a relationist solution to the earlier (abridged) version of Frege's puzzle. Recall that our problem was to locate a semantic difference between the pairs 'Cicero', 'Cicero' and 'Cicero', 'Tully' that did not require a semantic difference between the names 'Cicero' and 'Tully' themselves. But given our current notion of semantic fact, we can say what this difference is. For let us define two names to be *strictly* coreferential if it is a semantic fact that they corefer. The names 'Cicero' and 'Cicero' are then strictly coreferential while the names 'Cicero' and 'Tully' are not. Although the two pairs of names are both coreferential, it is a semantic requirement that they be coreferential in the one case but merely a semantic accident - if I may put it that way - in the other. Moreover, this semantic difference between the pairs is perfectly compatible, under a relationist view, with there being no intrinsic semantic difference between the names themselves.

We therefore have the beginnings of a relationist account of names. In the following lecture, I wish to develop the account further and to consider, in particular, its implications for the notion of content. This will then lead to a relationist account of thought and of the connection between thought and language.

