Lecture 1

INTRODUCTION

Many philosophers and linguists have remarked on the great expressive capacity of language - its capacity, on the basis of a finite vocabulary and a finite stock of semantic rules, to express an infinitude of different thoughts. But equally remarkable, though rarely remarked, is a capacity in the opposite direction - a capacity not to express different thoughts, but the very same thought from one occasion to the next. I say 'Cicero is an orator'; I say again 'Cicero is an orator'; you say 'Cicero is an orator'; you say again 'Cicero is an orator'; and the rest of mankind repeats what we say. Although we produce an infinitude of different utterances, we all somehow manage to express the very same thing. But how?

Perhaps the reason this contrasting capacity has gone unremarked is that it is not thought to be remarkable. After all, if I have said something once, then what is the point in saying it again? But such a response is completely off the mark. Just imagine that for some reason we were not able to say the same thing from one occasion to another. Reasoning would then be impossible, since it rests on a match in meaning or what is said. To take a simple illustration, the validity of modus ponens - the inference from sentence of the form 'S' and 'if S then T' to the sentence 'T' - depends upon there being no ambiguity in the use of the sentences 'S' or 'T' . Transmission of information would also be impossible. I inform you in the words 'Cicero is an orator' that Cicero is an orator. But you cannot pass the information on since you are unable to say what I said. It is not even clear that you can *believe* what I said - or, at the very least, you cannot both believe what I said and express what you believe, since that again would require that you say what I said. Nor is it clear that one could learn the use of language from others, since that requires the ability to use the words in the same way as the others. And how could I do that without being able to reproduce what you say?

There is another reason the contrasting capacity might have gone unremarked. For it might have been thought to be evident in what it consists. Language largely consists of the conventional association of meanings with expressions. As long as you and I associate the same meanings with the same expressions - or as long as I do from one occasion to the next - then we can say the same things by uttering the same words. Of course, there are complications arising from context, since what is said may vary from context to context even when the words are the same and their meanings are the same. But the context may be irrelevant or it may be relevant but stay the same. And even when it varies, we may adjust for the variation. If you say 'I am tired' then *I* can say 'you are tired'.

This obvious account of the capacity might be called the 'resemblance view'. According to this view, the capacity for reproduction depends upon the capacity for different utterances to resemble one another in respect of meaning - each means what it does and is thereby able to say the same thing.

Now there is a way in which the resemblance view might be completely innocuous. If you and I say the same thing, then our utterances must have something in common. For my utterance says the same as my utterance and your utterance also says the same as my utterance. They therefore have in common that each says the same as my utterance. But this is completely trivial: it does nothing to *account* for the fact that the two utterances say the same thing; it simply presupposes that they do.

If the resemblance view is to have any explanatory import, then it must be based upon the idea that the two utterances will have some *intrinsic* feature in common and that it is this which accounts for their saying the same thing. Two identical twins look alike; and they look alike in virtue of having some common intrinsic features. And it must be supposed that this is how it is with the utterances; they are semantic twins, as it were, that bear their semantic features on their face. Put it this way. We can take a ‘semantic snapshot’ of my utterance, one that reveals its semantic features without regard to its semantic relationship to your utterance. We can also take a semantic snapshot of your utterance, one that reveals its semantic features without regard to its semantic relationship to my utterance. The view is that it can then be determined on the basis of these semantic snapshots whether or not the two utterances say the same thing.

The resemblance view is very attractive and it is one that has been shared, as far as I know, by every major approach to the theory of meaning. Of course, part of what might make the view seem so attractive is the easy slide from the innocuous version of the resemblance view to the more substantive version. But far more significant is the fact that it is hard to see either why an alternative is necessary or what it might be. Two utterances that say the same thing clearly do have some common intrinsic semantic features; and it is not at all clear either why these are not sufficient to guarantee that they say the same thing or, if they do not, what the additional factors that guarantee that they do might be.

All the same, I wish to argue that the view is mistaken - and deeply mistaken at that. What I would like to propose in its place is a relational view of meaning. According to this view - which I call ‘semantic relationism’ - the fact that two utterances say the same thing is not entirely a matter of their intrinsic semantic features; it may also turn on semantic relationships among the utterances or their parts which are not reducible to those feature. Thus we must recognize that there may be *irreducible* semantic relationships, ones not reducible to the intrinsic semantic features of the expressions between which they hold. Thus, even if we take a semantic snapshot of each expression in our language, one that completely displays its intrinsic semantic features, it may not be evident from these snapshots what semantic relationships should hold among those expressions. The picture of an assemblage of semantic snapshots must therefore be supplemented by a picture in which these snapshots are connected, one to the other, by semantic threads. This snapshot goes with this one, that one with that; but there is no determining from the snapshots themselves how they are to be threaded.

It is important to distinguish semantic relationism from the more familiar doctrine of semantic holism. The underlying dispute between the holists and their opponents is over the proper form of semantic explanation. For the holists, a proper account will be broadly inferential in character - it will deal with the inferential relationships among different expressions or with other aspects of their inferential or conceptual role; while for their opponents, it will be representational in character - it will attempt to describe the relationship between the expressions of the language and the reality it is meant to represent. But both sides to the dispute will agree that if a representational semantics is adopted then it should be atomistic in form while if an inferential semantics is adopted it should be holistic in form.

It is, however, on this first point of agreement that I wish to situate the current dispute. Our concern is to reject Atomism - or what, in this context, I prefer to call 'Intrinsicalism' - for a representational form of semantics; and so our disagreement is as much with the holists as with their opponents.

We can better understand what is distinctive about the present position - and also how radical it is - by means of an analogy with the substantivalist conception of space. The substantivalist believes that a fundamenatal account of the spatial facts should incorporate an assignment of location to each spatial object; and similarly, the representationalist will think that a fundamental account of the semantic facts should incorporate an assignment of meaning, or representational role, to each meaningful expression. Now it would be quite bizarre for the substantivalist to believe that the fundamental spatial facts should also include spatial relationships among the spatial objects. For the natural - almost irresistible - view is that the spatial relationships among objects will be induced by the corresponding spatial relationships among their locations. What is it for two objects to be coincident? It is for their locations to be the same. What is it for two objects to be a foot apart? It is for their locations to be a foot apart. And similarly for all other spatial relations. A corresponding view about semantical relationships would appear to be equally plausible for the representationalist. What is it for two expresssions to be synonymous? It is for their meaning to be the same. What is it for them to be contraries? It is for their meanings to be contraries. And similarly for all other semantical relationships. But it is exactly this analogous view that I wish to deny. Not all semantical relationships among expressions are induced by corresponding relationships among their meanings and even synonymy - which might be thought to be a paradigm of an induced relationship - is not always properly so regarded.

These lectures will be organized around certain puzzles concerning what one might call the ‘reproduction' of of meaning. There are four main puzzles in all. The first, which is not well-known but which I believe can be traced back to Russell’s ‘Principles of Mathematics’, may be called the ‘antinomy of the variable’. The others are: Frege’s puzzle over names; Kripke’s puzzle concerning belief reports; and Moore’s paradox of analysis (which is also related to Mates’ puzzle over synonyms). It is my view that none of these puzzles can be adequately solved without appeal to relational ideas; and, indeed, as we progress through the puzzles we shall find ourselves committed to more and more far-reaching forms of relationism. Very roughly, the antinomy of the variable will show that relationism is true of variables, Frege's and Kripke's puzzles will show that it is true of names and singular thoughts, and Moore's paradox will show that it is true of predicates and concepts. Thus it will be seen that the semantic, or representational, threads run through the whole of language and thought and not just through this or that part.

It is generally supposed - by logicians and philosophers alike - that we now have a perfectly good understanding of how variables work in the symbolism of logic and mathematics. Once Frege had provided a clear syntactic account of variables and once Tarski had supplemented this with a rigorous semantic account, it would appear that there was nothing more of any significance to be said. It seems to me, however, that this common view is mistaken. There are deep problems concerning the role of variables that have never been properly recognized, let alone solved, and once we attempt to solve them we shall see that they have profound implications not only for our understanding of variables but also for our understanding of other forms of expression and for the general nature of semantics.

At the root of the present difficulties is a certain puzzle. In order to state this puzzle, we shall need to appeal to the notion of *semantic role*. I do not mean this as a technical notion of a kind that one might find in a formal semantics, but as a non-technical notion whose application may already be taken to be implicit in our understanding of a given language or symbolism. For in any meaningful expression, there is something conventional about the expression - having to do with the actual symbols or words used, and something non-conventional - having to do with the linguistic function of those symbols or words. And 'semantic role' is just my term for this linguistic or non-conventional aspect of a meaningful expression.

Now for the puzzle. Suppose that we have two variables, say 'x' and 'y'; and suppose that they range over the same domain of individuals, say the domain of all reals. Then we would like to say the following about their semantic role. First:

Semantic Sameness (SS): The semantic roles of the variables x and y are the same.

This is evidenced by the fact that the difference between the expressions 'x > 0' and 'y > 0', say, is purely notational; it lies merely in the choice of the symbol and not in linguistic function. Second:

Semantic Difference (SD): The semantic relationships between the pairs of variables x, x and x, y are different.

This is evidenced by the fact that there is a clear difference in semantic role between 'x > x' and 'x > y'; for the first is true for no values of its variables while the second is true for some values of its variables and not true for others. And how could there be a difference in semantic role between these two formulas without there being a corresponding difference in the semantic relationships between the pairs of variables that they involve?

Third:

Intrinsicality (I) If there is a difference in semantic relationship between the pairs x, x and the pairs x, y then the semantic role of the variables x and y is not the same.

For clearly x has the same semantic role as itself. And if x also had the same semantic role as y, then how could x, x fail to have a different semantic role from x, y?

However, the three assumptions are inconsistent; and so one must be rejected. But which?

One naturally supposes that the solution to our puzzle should be sought in the various semantics that have been developed for the language of predicate logic. After all, it is presumably the aim of these semantics to account for the semantic role of the expressions with which they deal; and so we should expect them to account, in particular, for the semantic role of variables.

However, when we turn to the various semantics, we find them entirely unsuited to the purpose. Let us begin with the most familiar of them, that of Tarski [36]. The reader will recall that the Tarski semantics proceeds by defining a relation of satisfaction between assignments and formulas. To fix our ideas, let us suppose that the variables of our language are x_1, x_2, ... and that the domain of discourse is *D*. We may then take an *assignment* to be a function taking each variable of the language into an individual from *D*; and the semantics will specify - by means either of a definition or of a set of axioms - what it is for each kind of formula to be satisfied by an assignment.

Now what account, within the framework of the Tarski semantics, can be given of the semantic role of the variables? There would appear to be only two options. The first is to take the semantic role of a variable to be given by its range of values (the domain *D* in the case above). Indeed, quite apart from the connection with the Tarski semantics, this is the usual way of indicating how a variable is to be interpreted: one specifies its range of values; and that is it.

This approach does indeed account for the fact that the semantic role of any two variables x and y (with an identical range of values) is the same. But it does nothing to account for the semantic difference between the pairs of variables x, y and x, x; and nor is any reason given for denying that there is an intuitive difference in semantic role.

The other option is to take the semantic role of a variable to be what one might call its 'semantic value' under the given semantics. The semantic values are those entities which are assigned (or which might be taken to be assigned) to the meaningful expressions of the language and with respect to which the semantics for the language is compositional. When we examine the Tarski semantics, we see that the semantic value of an open formula (one containing free variables) might be taken to be the function that takes each assignment into the 'truth-value' of the formula under that assignment and, similarly, the semantic value of an open term might be taken to be the function that takes each assignment into the denotation of the term under that assignment. We then easily see that the Tarski semantics is compositional with respect to the semantic values as so conceived; it 'computes' the semantic value of a complex expression on the basis of the simpler expressions from which it is derived.

Under this conception of semantic value, the semantic value of a variable x will be a function which takes each assignment into the individual which it assigns to x. It is therefore clear, if we identify semantic roles and semantic values, that x and y will differ in their semantic roles; for if we take any assignment which assigns different individuals to x and y, then the semantic value of x will deliver the one individual in application to that assignment while the semantic value of y will deliver the other individual in application to the assignment.

We therefore secure the semantic difference between the pairs x, y and x, x under this account of semantic role. But we are at a loss to account for the fact that the semantic role of the variables x and y is the same; and nor is any reason given for disputing the intuitive identity in semantic role. Tarski's semantics does not therefore provide the materials for solving the antinomy of the variable.

There are, however, other semantics on the market. What of them? The two main candidates are what I call 'the instantial' and the 'algebraic' accounts. According to the instantial account, the semantic value of a quantified sentence such as $\exists x(x > 0)$ is made to depend upon the semantic value of a closed instance $c > 0$ (e.g., $3 > 0$). The intuitive idea behind the proposal is that, given an understanding of a closed instance $c > 0$, we thereby understand what it is for an arbitrary individual to satisfy the condition of being greater than 0 and are thereby in a position to understand what it is for some individual or other to satisfy this condition.

According to the second account, the semantic value of a quantified sentence such as $\exists x(x > 0)$ is made, in the first place, to depend upon the semantic value of the corresponding λ-term $\lambda x(x > 0)$, denoting the property of being greater than 0. Of course, this merely pushes the problem back a step, since we now need to account for the semantic value of the λ-terms. But this may be done by successively reducing the complexity of the λ-terms. The semantic value of $\lambda x{\sim}(x > 0)$, for example, may be taken to be the 'negation' of the semantic value of $\lambda x(x > 0)$, while the semantic value of $\lambda xy(x \leq y \vee y \leq x))$ may be taken to be the 'disjunction' of the semantic values of $\lambda xy(x \leq y)$ and $\lambda xy(y \leq x)$. In this way, the λ- bindings may be driven inwards to the atomic formulas of the symbolism and their application to the atomic formulas may then be replaced by the application of various 'algebraic' operations to the properties or relations signified by their primitive predicates. $\lambda x(x > x)$, for example, might be taken to be the result of reflexivizing the greater-than relation.

These two accounts have in common that they do not assign a semantic value to open terms or formulas; at each stage in the process of semantic evaluation, we deal only with closed expressions. This makes their relationship to the antinomy somewhat unclear. Under a weak interpretation, the accounts are taken to be silent on the question of the semantic role of free variables. But this means that they do nothing to solve the antinomy, though they may do something to undermine its significance since they purport to show that, even if free variables have a semantic role, there is no need to appeal to that role in providing an semantic account of quantification or other forms of variable-binding. Under a strong interpretation, the accounts are taken implicitly to deny that free variables have any semantic role. They are syncategorematic, like brackets or punctuation; and the antinomy is based upon the false presupposition that variables have a semantic role.

There is a great deal to be said about the strong interpretation of these accounts. But let me merely remark that the intuitive evidence against the view appears to be overwhelming. Surely it is at least part of the semantic role of an open term to represent a range of values. It is part of the semantic role of the term ‘2n’, for example, to represent any even number and part of the semantic role of the term ‘ai + b’ to represent any complex number. Just as it is characteristic of a closed term such as ‘2.3' to represent a particular individual, so it is characteristic of an open term, such as ‘2n’, to represent a range of different individuals; and just as the representation of a particular individual is a semantic relationship, so is the representation of a range of individuals. We would therefore appear to have as much a reason to regard the representation of a range of individuals as a part of the semantic role of an open term as we have to regard the representation of a particular individual as part of the semantic role of a closed term. The intuitive evidence therefore suggests that the antinomy is still with us.

But even if we put this evidence aside, I think it can be shown that these accounts are inadequate in their own terms. Let us briefly consider each in turn. According to the instantial account, it will be recalled, a closed quantified sentence - such as $\exists x B(x)$ - is to be understood on the basis of any one of its instances B(c). Now we need not deny that if we understand the closed instance B(c), then we are also capable of understanding the quantified statement (given, of course, that we understand the apparatus of quantification). But it is a mistake to suppose that our understanding of the quantified sentence is *derived* from our understanding of a particular instance since there may be no particular instance that we are in a position to understand. Suppose, for example, that the variables range over all points in abstract Euclidean space. Then it is impossible to name any particular point. But if we are capable of understanding any instance of the quantified sentence then, a fortiori, we are incapable of deriving our understanding of the quantified sentence from an understanding of an instance.

It would be more accurate to say that our understanding of a quantified sentence is derived, not from our understanding of a *particular* instance, but from our understanding of an *arbitrary* instance. We have a general understanding of what the semantic value of an instance might be; and it is from this general understanding that we derive our understanding of the quantified sentence. But this general understanding is more naturally taken to be implicit in our understanding of the open formula B(x) rather than in our understanding of any particular instance.

There is another serious problem with the approach, for it cannot be accommodated within an extensional framework. Philosophers who advocate the instantial approach never bother to make clear how it might be made precise. What exactly is the semantic value of the closed instance? And how exactly is the semantic value of the corresponding quantified sentence to be derived? But once we give the matter some thought, it becomes clear that the semantic value of the instance must be an intensional entity that shares, to a high degree, in the structure of the sentence to which it corresponds. For let us, without prejudice, call the semantic value assigned to the instance B(c) a ‘proposition’. Then in order to derive the semantic value of the corresponding quantified sentence $\exists xB(x)$, we must make sense of‘removing’ the entity corresponding to the constant c from this proposition in order to obtain a propositional form to which the quantifier concept might then be applied. But how are we to make sense of this operation unless we presuppose that the proposition is already endowed with some sentence-like structure?

[So much the worse, one might think, for the extensional approach. But however sympathetic one might be to alternative semantic approaches, it is hard to believe that our current problems lie in the adoption of an extensional approach. After all, the extensionalist credentials of variables are as good as they get: they simply range over a given domain of individuals without the intervention of different senses for different individuals and without the need for different senses by which the domain might be picked out for different variables. It is therefore hard to see why the addition of variables to a language that was otherwise in conformity with extensional principles should give rise to any essential difficulties. If the extensional project fails, it cannot be because the variables carry some hidden intensional freight.]

[I might also note, as an historical aside, that the views of Frege [52] on the sense/reference distinction require that it be possible to provide an extensional semantics for the language of predicate logic, since he wanted to be able to provide a compositional semantics at the level of both sense and reference. The instantial approach is sometimes attributed to Frege (e.g., by Dummett [73], pp. 15-16). But it is only intelligible, at best, at the level of sense. And, for Frege, it is not even a viable option at the level of sense, since his views require that there be a *parallel* semantics at the level of reference.]

In certain ways, the algebraic account fares even worse. In the first place, there are difficulties in seeing how it is to be extended beyond the symbolism of predicate logic. Here is one example, though there are others. Suppose that we add an operator ‘$\Box$’ for necessity to the symbolism for first-order logic and that we take the quantifiers to be ‘actualist’ - ranging, in each possible world, over the objects that exist in that world. Consider now the algebraist’s treatment of $\Box(x = x)$ and of $\Box\exists y(y = x)$. In conformity with the algebraic approach, $\lambda x\, \Box(x = x)$ should be understood to signify the result of applying some operation, call it ‘necessitation’, to the property (of self-identity) signified by $\lambda x(x = x)$ and $\lambda x\Box\exists y(y = x)$ should be understood to signify the result of applying this same operation to the property (of existence) signified $\lambda x\exists y(y = x)$. But under an actualist interpretation of the quantifier (and hence also of λ-binding), $\lambda x(x = x)$ and $\lambda x\exists y(y = x)$ will be modally indistinguishable - they will be true, in each possible world, of the individuals that exist in that world. And this means that $\Box(x = x)$ and of $\Box\exists y(y = x)$ will be modally indistinguishable - which, of course, they are not, since the first is satisfied by any individual whatever while the second is satisfied only by individuals that necessarily exist.

Thegeneral point is that the success of the algebraic approach depends upon presupposing the truth of certain ‘distribution’ principles and there is no reason, in general, to suppose that such principles will be true. But one cannot accept as satisfactory an account of quantification that depends upon certain special features of the language to which it is applied.

Another serious difficulty with the approach is that it requires us, at almost every turn, to make arbitrary decisions about the interpretation of the symbolism, decisions which have no counterpart in our actual understanding of the symbolism. Let me merely give one illustration. How are we to interpret the application of a λ-operator to an atomic formula? How, for example, are we to understand $\lambda zxRxxz$? We could take this to be the result of first forming the ‘converse’ $\lambda zxyFxyz$ of $\lambda xyzFxyz$ and then forming the reflexive version $\lambda zxFxxz$ of the converse, or we could take it to be the result of first forming the reflexive version $\lambda xzFxxz$ of λxyzFxyz and then forming the converse λzxFxxz of the reflexive version; and similarly, and to a much greater degree, for other cases.

These choices do not correspond to anything in our actual understanding of the symbolism; and the need to make them suggests that the approach basically misconceives in what that understanding consists.

I now wish to indicate how I think the antinomy is to be solved and how a more satisfactory semantics for the symbolism of first-order logic might be developed on the basis of its solution.

We accept the sameness and difference principles and, in keeping with our relationist credentials, we reject the connecting principle that renders them inconsistent.

This is likely to go against the grain. How, it might be asked, how can there be a difference in the semantic relationships holding between each of two pairs of expressions without there being an intrinsic difference in the semantic features of the expressions themselves? Thus given a difference in semantic relationship between the pairs 'doctor', 'dentist' and 'doctor', 'doctor' (with the one being synonymous and the other not), then surely must be a difference in meaning between 'doctor' and 'dentist'. And similarly, given a difference in semantic relationship between the pairs x, y and x, x, then surely there must be a semantic difference in role between the variables x and y.

It has to be acknowledged that this view of meaning - I call it 'semantical intrinsicalism' - is extremely plausible. But it has to be given up. For once one examines the actual behavior of variables, without theoretical preconception, then it can clearly be seen to be in conflict with the view. The case of variables constitutes a clear counter-example.

Suppose, as before, that we are dealing with a language that contains the variables x_1, x_2, x_3, ... How then is their semantic behavior to be described? For simplicity, let us suppose that their semantic behavior is to be described entirely in extensional terms, since nothing in our argument will turn upon allowing intensional elements to appear in the description.

Now we should certainly specify the range of values each variable can assume and, of course, as long as the language is 'one-sorted', the range of values for each variable will be the same. Now normally, in providing some kind of semantical description of the variables, nothing more is said (perhaps because of the grip of intrinsicalist doctrine). But something more does need to be said. For we should specify not only which values each *single* variable can assume, when taken on its own, but also which values *several* variables can assume, when taken together. We should specify, for example, not only that x_1 can assume the number 2 as a value, say, and x_2 the number 3 but also that x_1 and x_2 can *simultaneously* assume the numbers 2 and 3 as values; and, in general, we should state that the variables take their values independently of one another, that a variable can take any value from its range regardless of which values the other variables might assume.

It is here important to appreciate that it does not follow, simply from the specification of a range of values for each variable, which values the variables can *simultaneously* assume. One might adopt the proposal of Wittgenstein [1922], for example, and disallow distinct variables from taking the same value; or, at the other extreme, one might insist that distinct variables should always assume the same value (treating them, in effect, as strict notational variants of one another); and there are, of course, numerous other possibilities. Thus the fact that distinct variables assume values in complete independence of one another is an additional piece of information concerning their semantic behavior, one not already implicit in the specification of their range.

However, once we have specified the range and the independence in value, then we will have a complete description of the semantic behavior of the variables; there is nothing more (at least at the extensional level) to be said about their role. But if this is so, then it is clear that the intrinsicalist doctrine, *no difference in semantic relationship without a difference in semantic feature*, will fail. For the intrinsic semantic features of any two variables will be the same - it will in effect be given by the specification of their range, whereas the intrinsic semantic features of the pairs x_1, x_2, say, and x_1, x_1 will be different, since the former will assume *any* pair of values from the given range while the latter will only assume *identical* pairs of values. If we are merely informed of the intrinsic semantic features of two variables, then we cannot tell whether they assume their values independently of one another (should they be distinct) or whether they always assume the same value (should they be same).

It is thus by giving up the intrinsicalist doctrine, plausible as it initially appears to be, that the antinomy is to be solved.

Let us now consider how the relational solution to the antinomy is capable of yielding a more satisfactory semantics for the symbolism of first-order logic.

I take it to be the aim of a semantics for a given language to account for the semantic behavior of its expressions. This is usually done is by assigning a semantic value to each meaningful expression of the language, with the semantic value of a complex expression being determined on the basis of the semantic values of the simpler expressions from which it is syntactically derived. The semantic value of each expression is naturally taken to correspond to its intrinsic semantic role; and so, given the truth of semantic intrinsicalism, the assignment of semantic values to the expressions of a language should then be sufficient to determine their semantic behavior.

But we have seen that the doctrine of semantic intrinsicalism should be abandoned; there are (intrinsic) semantic relationships between expressions that are not grounded in their intrinsic semantic role. This means that the aim of semantics should be reconsidered. For it will no longer serve to assign semantic values to expressions; we should also take account of semantic relationships between expressions that may not be grounded in their intrinsic semantic features. Let us use ‘semantic connection’ as the formal counterpart, within a semantics, to the informal notion of semantic relationship (just as ‘semantic value’ is the formal counterpart, within a semantics, to the informal notion of semantic role). Then semantic connection should replace semantic value as the principal object of semantic enquiry and the aim of semantics should be the determination of semantic connections among expressions on the basis of the semantic connections among simpler expressions. (Semantic values merely represent the special case of the degenerate connection on a single expression.)

Following through this strategy, the semantics for a given language will eventually terminate in the ‘lexical semantics’, which accounts for the behavior of those expressions, the ‘lexical items’, that are not derived from any other expressions. However, the lexical semantics, like the semantics as whole, must now be taken to assign semantic connections to the lexical items, and not merely semantic values. It is on the basis of these ‘primitive’ semantic connections that the semantic connections among all expressions of the language will ultimately be determined. This is the general idea of a relational semantics.

If we are to apply this general idea to the language of first-order logic, then we must first have some conception of what we want the semantic connections to be. 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. Similarly, given a sequence of expressions, we may takes its value-range - or semantic connection - to be the set of sequences of values that the expressions are simultaneously capable of assuming. So, for example, the semantic connection on 'x + y', 'x > y', 'z' will include the sequences 5, FALSE, 6 and 6, TRUE, 2. It should be noted that the semantic connections are entirely non-typographic; they contain no trace of language. There is therefore no danger of our semantics being implicitly typographic.

We must now show how to determine the semantic connection on any given sequence of expressions - starting with the lexical semantics, for the very simplest expressions, and then successively working through more and more complicated forms of expression. 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. It is at this point in the semantics that relationism enters the scene.

We now need rules for extending the semantic connections to more complicated expressions and more complicated sequences of expressions. Let us 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. But 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 then 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 give us the result we want.

Quantifiers raise some additional problems, which we should briefly mention. 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 $x > 0$, x (or, more explicitly, on $\exists$, x, $x > 0$, x). But now the two occurrences of x that derive from the quantified formula will become accidentally coordinated or 'bound' with the other occurrence of x.

In order to solve this difficulty, we must allow ourselves to relax the assumption that distinct occurrences of the same free variable always take the same value. Even though the variables are the same, they may behave as though they are different. This means that in the symbolism we use in presenting the semantics on sequences we must explicitly represent when two occurrences of the same variable are or are not to be coordinated. Thus any sequence of expressions will be equipped with a 'coordination scheme' which explicitly indicates the intended pattern of coordination on the free occurrences of variables. This is an extremely important idea; and it is one to which we shall return. But let us for now merely note that it provides us with a straightforward solution to the problem over quantification for, when we evaluate $\exists x(x > 0)$, x in terms of x, $x > 0$, x, we should require that the first two occurrences should be coordinated with each other but not with the third.

The relational semantics has several clear advantages over its rivals. First and foremost, it embodies a solution to the antinomy: the intrinsic semantic features of x and y (as given by the degenerate semantic connections on those variables) are the same, though the intrinsic semantic features of the pairs x, y and x, x (again as given by the semantic connections) are different. But the semantics is also more satisfactory, in various ways, *as* a semantics. In contrast to the algebraic and instantial approaches, it accounts for the semantic role of open expressions; in contrast to the instantial approach, it can be provided with an extensional formulation; and in contrast to the algebraic approach, it is based upon a credible method of evaluation. Finally, in contrast to the Tarski semantics, it is non-typographic. By going relational, we avoid having to incorporate the variables themselves into the very identity of the entities that the semantics assigns to the meaningful expressions of the language.