The Role of Variables

It is generally supposed - by logicians and philosophers alike - that we now possess 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 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 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.

It is my aim in the present lecture to explain what these problems are and how they are to be solved. I begin with an antimony concerning the role of variables which I believe any satisfactory account of our understanding of them should solve (§1). I then argue that the three main semantical schemes currently on the market - the Tarskian, the instantial and the algebraic - are unsuccessful in solving the puzzle (§2-3) or in providing a satisfactory semantics for first-order logic (§4-5). Finally, I offer an alternative scheme that it is capable of solving the antimony (§6) and of providing a more satisfactory semantics for first-order logic (§7). It is based upon a new approach to representational semantics, which I call semantic relationism ; and in the remaining three lectures, I will discuss the implications of this approach for the semantics of names and belief-reports.

§1. The Antinomy of the Variable

At the root of the present considerations 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 intuitive 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.

In fact, there will be no need, in stating the puzzle, to appeal to semantic role per se. For our interest will solely be in whether two expressions have the *same* or a *different* semantic role. And, of course, their having the same, or different, semantic roles is compatible with different views as to what that semantic role might be. Indeed, our puzzle would appear to arise under any reasonable conception of semantic role - be it Russellian, Fregean, or of some other kind.

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 it appears as if we wish to say contradictory things about their semantic role. For when we consider their semantic role in two distinct expressions - such as x > 0' and y > 0', we wish to say that their semantic role is the same. Indeed, this would appear to be as clear a case as one could hope to have of a mere conventional or notational difference; the difference is merely in the choice of the symbol and not in its linguistic function. On the one hand, when we consider the semantic role of the variables in the same expression - such as $x > y$ - then it seems equally clear that their semantic role is different. Indeed, it is essential to the linguistic function of the expression as a whole that it contains two distinct variables, not two occurrences of the same variable, and presumably this is because the roles of the distinct variables are not the same.

Generalizing from our example, we arrive at the following two claims:

Semantic Sameness (SS): Any two variables (ranging over a given domain of objects) have the same semantic role; and

Semantic Difference (SD): Any two variables (ranging over a given domain of objects) have a different semantic role.

Yet we cannot univocally maintain both (at least given that there are two variables!).

We might call this puzzle the *antimony of the variable*. It was first stated by Russell, though in ontological rather than semantical fashion. He writes ([1903], §93), x is, in some sense, the object denoted by *any term*; yet this can hardly be strictly maintained, for different variables may occur in a proposition, yet the object denoted by *any term*, one would suppose, is unique. It also bears a close affinity with Frege s puzzle concerning names (Frege [1952]), though with variables taking the place of names.

Clearly, any resolution of the conflict must begin by locating an ambiguity in the phrase same semantic role . Whatever the sense in which the variables x and y are semantically the same in $x > 0$' and $y > 0$' cannot be the sense in which they fail to be semantically the same in $x > y$. Now as a first stab towards locating the ambiguity, we might appeal to a notion of context. The variable x in the context of the formula $x > 0$' plays the same semantic role as the variable y in the context of the formula $y > 0$'. On the other hand, the two variables x and y will play different semantic roles within the context of the single formula $x > y$. Thus (SD) will hold for sameness of semantic role in the *cross*-contextual sense while (SS) will hold for difference of semantic role in the *intra*-contextual sense; and, given the ambiguity in the respective senses of same and different role, contradiction is avoided.

Natural as this response may be, it does not really solve the puzzle but merely pushes its solution back a step. For why do we say that the variables x and y have a different semantic role in $x > y$? Clearly, it has to do with the fact that the occurrence of y cannot be replaced by x without a change in semantic role; the role of x, y in $x < y$ is different from the role of x, x in $x < x$. In other words, the intra-contextual difference in semantic role between the variables x and y within the single formula $x > y$ amounts to a cross-contextual difference in semantic role between the pair of variables x, y in $x > y$ and the pair of variables x, x in $x > x$. And, in general, to say that there is an intra-contextual difference between x and y, in the intended sense, is just to say that there is cross-contextual difference between the pair of variables x, y and the pair x, x.

We may therefore state (SS) and (SD) in the following forms:

(SS)′ there is no cross-contextual difference in semantic role between the variables x and y; and

(SD)′ there is a cross-contextual difference in semantic role between the pair of variables x, y and the pair x, x,

using an univocal notion of semantic role throughout.

In contrast to the earlier formulation, there is now no explicit contradiction. But there is still a difficulty in seeing how the two claims, (SS)′ and (SD)′, might both be true. For how can there be an cross-contextual difference in semantic role between the pair of variables x, y and the pair x, x unless there is an cross-contextual semantic difference between the variables x and y themselves? What else could account for the difference in semantic role between the pairs x, y and x, x except a semantic a difference in the individual variables x and y? Or to put it another way, if there is a semantic difference between x, y and x, x, then there must be a semantic difference between x and y; and it is hard to see why this difference should only be turned on or made manifest when the variables are used in the same context and not when they are used in different contexts.

The puzzle therefore remains; and any solution to the puzzle should either explain how (SS)′ and (SD)′ might be compatible, notwithstanding appearances to the contrary, or it must explain how one of (SS)′ or (SD)′ might reasonably be rejected.

§2. The Tarskian Approach

It might be thought 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 that have in fact been developed, 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. It will state for example that an assignment satisfies the formula ±B just in case it fails to satisfy B or that an assignment satisfies the formula ±xB just in case every x-variant of the assignment (differing at most on the value assigned to x) satisfies B.

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 (the TRUE if the formula is satisfied by the assignment and the FALSE otherwise) 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. Thus the semantic value of the formula x > 0 (under the natural interpretation of the language) would be the function that takes any assignment into TRUE if the number it assigns to x is positive and into FALSE otherwise; and the semantic value of the term x + y would be the function which takes an assignment into the sum of the numbers which it assigns to x and y. 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 special case of the semantic value of a term; it 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 (this requires, of course, that there be at least two individuals in the domain), 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 unable 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 of semantic role.

There is a another, perhaps more serious, problem with the approach. For although it posits a difference between the variables x and y (and hence between the pairs x, y and x, x), it does nothing to account for their *semantic* difference. For in the last analysis, the posited difference between the semantic values for x and y simply turns on the difference between the variables x and y themselves. The semantic value of the variable x, for example, delivers one individual rather than another in application to a given assignment simply because it is the variable that it is. Thus what we secure, strictly speaking, is not a *semantic* difference, one lying on the non-conventional side of language, but a *typographic* difference, one lying purely on the conventional side of language; and so we have done nothing, properly speaking, to *account* for the semantic difference between x and y (or between x, y and x, x).

§3 The Denial of Semantic Role

In stating the antimony of the variable, we have presupposed that variables *have* a semantic role; and it might be thought that this is the root cause of our difficulties. For it might be thought that our understanding of variables is inseparable from their role in quantification or other forms of variable-binding and that any attempt to explain the role of *free* variables, apart from their connection with the apparatus of binding, must therefore fail.

There is a familiar approach to the semantics of predicate logic that appears to lend some support to this point of view. For in attempting to provide a semantics for quantified sentences, we face a problem that is in some ways analogous to our antimony. We wish to assign a semantic value to a quantified sentence, such as $\exists x(x > 0)$; and we naturally do this on the basis of the semantic value assigned to the open sentence $x > 0$ that is governed by the quantifier. But, given that $x > 0$ and $y > 0$ are mere notational variants, they should be assigned the same semantic value; and so we should assign the same semantic value to $\exists x(x > 0)$ and $\exists x(y > 0)$ - which is clearly unacceptable.

Now one solution to this problem is to deny that the semantic value of $\exists x(x > 0)$ is to be assigned on the basis of the semantic value assigned to $x > 0$; and once this line is adopted, then consistency demands that we never appeal to the semantic value of an open expression in determining the semantic value of a closed expression. In other words, the semantics for closed expressions should be autonomous in the sense of never making a detour through the semantics of open expressions.

There are two main ways in which autonomy of this sort might be achieved. I call them the instantial and the algebraic approaches respectively. According to the first, 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 this 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.[1]

According to the second approach, 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 $\lambda x\neg(x > 0)$, for example, may be taken to be the negation of the semantic value of $\lambda x(x > 0)$, while the semantic value of $\lambda xy(x < y \vee y < x)$ may be taken to be the disjunction of the semantic values of $\lambda xy(x < y)$ and $\lambda xy(y < x)$. 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 the primitive predicates.[2]

In discussing these proposals, it is important to distinguish between two different questions. The first is whether they are plausible or even viable. Can a semantics of the proposed sort be given and, if it can, then is it faithful to the way we actually understand the symbolism? The second question is whether open expressions should be taken to have a semantic role (which it might then be part of the aim of semantics to capture).

Of course, if there is an autonomous semantics for closed expressions, then that deprives us of one reason for thinking that open expressions have a semantic role, since they are not required to have a semantic role in order to account for the semantics of closed expressions; and I

---

[1] Cf Dummett [73], pp. 15-16.

[2] Bealer ([79], [82], [83]) is the leading contemporary advocate of this approach.

suspect that many philosophers who have been attracted to the idea of an autonomous semantics for closed expressions have been inclined, on this basis, to reject a semantic role for open expressions. It seems to me, however, that there are strong independent reasons for thinking that open expressions do indeed have a semantic role.

The intuitive evidence for this appears to be overwhelming. Surely, we are inclined to think, 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 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.

But the opponent of semantic roles for open expressions is unlikely to be impressed with these considerations. For he may argue that, in so far as we think of an open expression as having a semantic role, it is because we think of its variables as being implicitly bound. There is perhaps no need to think of them as being bound by a *quantifier*, thus there is no need to think of us as always using these expressions to express thoughts, but there must at least be some variable- binding operator in the background by means of which the supposed semantic role of the open expression is to be understood. So, for example, the term 2n might be understood, in so far as it is seen to have a semantic role, as doing duty for the set-term {2n: n a natural number} or for the λ-term λn.2n (denoting the function from each number to its double).

I do not regard this account of the alleged semantic role of open expressions as at all plausible. The alleged semantic role of open expressions, in terms of its representing a range of values, would appear to be perfectly intelligible quite apart from the possible connection with variable-binding. Indeed, some philosophers have supposed that certain types of variable - or variable-like expression - might not even be subject to quantification or other forms of variable-binding. The schematic letters A and B of Quine ([52], §1.5), for example, are meant to stand in for the sentences of some language and yet are not bindable on pain of supposing that the sentences are names for some special kind of entity. These schematic letters, as much as regular variables, are subject to the antimony and yet would appear to have an independent semantic role. Of course, Quine might be mistaken in his reasons for thinking that schematic letters are not bindable, but he is surely not mistaken in thinking that their semantic role can be understood apart from the connection with quantification or other forms of variable binding.

The opponent of semantic roles for open expressions also faces the awkward issue of saying what the implicit binding should be taken to be. One wants to say that a term 2n indifferently represents all even numbers. Our opponent says that the term can only be regarded as having a semantic role in so far as it is implicitly bound. But by what? Two obvious candidates, when no sentential context is at hand, are the set-term {2n: n a natural number} and the λ-term λn.2n ; a less obvious candidate is the *arbitrary even number* of Fine ([84], [85]). Thus it must be supposed that an implicit reference and an implicit ontological commitment is made to the set or function or arbitrary object, or to something else of this sort. But this is both arbitrary and gratuitous. For there is no reason to suppose that the implicit reference is to one of these entities as opposed to the other and it appears entirely irrelevant to our use of the term that it should carry any such implicit reference or commitment.

A perhaps even more decisive objection to the position arises from the consideration of semantic relationships. Not only do open expressions appear to have semantic roles, they also appear to enter into semantic relationships. For example, the value of the term $n + 1$ is always greater than the value of n, though not of m. But how is our opponent to account for these apparent semantic relationships? If he takes the variables of each term, taken on its own, to be implicitly bound, he is sunk: for then $n + 1$ will signify the successor function, say, and n the identity function, and so we will lose the special semantic relationship that holds between $n + 1$ and n as opposed to $n + 1$ and m. He must therefore take the variables of the two terms to be somehow simultaneously bound. He must say something like: what accounts for the apparent semantic relationship between $n + 1$ and n is the fact that the quantified sentence $\pm n(n + 1 > n)$ is true. But it seems bizarre to suppose that one must create this artificial context in which both terms occur in order to explain the semantic relationship between them. What kind of strange semantic relationship between the terms is it that can only be explained by embedding them within a richer language? Indeed, the proposed explanation of the semantic relationship presupposes that the relevant semantic features of the terms are preserved when they are embedded in the context of a single sentence; and so unless we had some independent way of explaining what that semantic relationship was, we would have no way to say what the presupposition was or whether it was correct.

If we are right, then the independent semantic role of open expressions is not to be denied and the antimony is not to be solved by denying that they have such a role.

§4. The Instantial Approach

I have so far left open the question of whether there might be an autonomous semantics for closed expressions, one not taking a detour through open expressions. But it seems to me that no such semantics is viable - or, at least, plausible. That this is so is not important for solving the antimony, but it is important for understanding the significance of a solution. For if a semantics for quantification or for other forms of variable binding must proceed by providing a semantics for open expressions then it should provide, in particular, a semantical account of free variables; and so no such account can be considered satisfactory unless it shows how the antimony is to be solved. Thus the failure of existing semantics to solve the antimony is a symptom of their inadequacy.

As I have mentioned, there are two main forms of autonomous semantics, the instantial and the algebraic, though the reasons for thinking them unsatisfactory are somewhat different. Let us begin with the instantial approach, which I regard as a closer approximation to how a satisfactory semantics should proceed.

According to this approach, it will be recalled, a closed quantified sentence, such as $\pm$B(x) is to be understood on the basis of one of its instances B(c) - the intuitive idea being that from an understanding of B(c), we may acquire an understanding of what it is for an arbitrary individual to satisfy the condition denoted by B( ) and that, from this, we may then acquire an understanding of what it is for this condition to be satisfied by some individual or other. But although the intuitive idea behind the proposal may be clear, it is far from clear how the proposal is to be made precise.

A certain semantic value is to be assigned to a closed instance B(c) of the existential sentence ±⊠B(x). Let us call it a proposition , though without any commitment to what it is. A certain condition is then to be determined on the basis of this proposition. But how? We took it to be the condition denoted by the scheme B( ) which results from removing all displayed occurrences of the term c from B(c). This suggests that the condition should likewise be taken to be the result of removing all corresponding occurrences of the individual denoted by c from the given proposition; indeed, we are given no other indication of how the condition might be determined. It must therefore be presupposed that there is an operation of abstraction which, in application to any proposition and any occurrences of an individual in that proposition, will result in a certain condition or propositional form in which the given occurrences of the individual have been removed. Once given such a form, we may then take quantified sentence ±⊠B(x) to predicate existence of it.

Now a great deal more needs to be said about the operation of abstraction before we have a precisely formulated semantics.[3] But one thing is clear. The use of such an operation in formulating the semantics for predicate logic is not compatible with an extensional approach, one in which we take cognizance only of the individuals denoted by the closed terms, the extensions of the predicates, and the truth-values of the sentences. For if there is to be a meaningful operation of abstraction, then the propositions to which it is to be applied must to some extent share in the structure of the sentences by which they are expressed; they must contain individuals in a way analogous to the way in which the sentences contain terms; and it must make sense to remove the individuals from the propositions in a way that is analogous to the removal of terms from a sentence. But clearly, there is nothing in an extensional approach that would enable us to make sense of such ideas.

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 adherence to extensionality. 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.

There is in any case another, more subtle, difficulty with the instantial approach which not even the intensional form of semantics is able to solve. For 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 incapable of understanding any

---

[3]Some hints at how a semantics of this sort might be developed are given in Fine [89], pp. 237-8. I hope to deal with the matter more fully elsewhere.

instance of the quantified sentence then, a fortiori, we are incapable of deriving our understanding of the quantified sentence from ourunderstanding of an instance.

Of course, what we really wanted to say was that the understanding of $\forall x B(x)$ should be derived from our understanding of an *arbitrary* instance. We have seen that this should not be taken to mean that our understanding of $\forall x B(x)$ derives from our understanding of some particular instance, though it does not matter which. But then what does it mean?

The only reasonable view seems to be that our understanding of $\forall x B(x)$ should be taken to derive from our general understanding of B(x), i.e. from our understanding of the proposition expressed by B(x) for any given value of the variable x. But the idea of a closed instance then falls by the wayside and we are left with the idea of an understanding the quantified sentence $\forall x B(x)$ in terms of the corresponding open sentence B(x). Thus we see that the instantial approach, once properly understood, does not even constitute an autonomous form of semantics.

§5. The Algebraic Approach

Under the alternative autonomous approach, the apparatus of binding is traded in for an algebra of operations, with the apparatus serving, in effect, as a device to indicate how the semantic value for a whole sentence is to be generated from the properties and relations expressed by the primitive predicates that it contains. In this case, there is no difficulty in making the semantics precise or in presenting it in extensional form. But there is a difficulty in seeing how to extend it beyond the standard symbolism of predicate logic.

One difficulty of this sort arises from the use of quantifiers that apply to several variables at once, though in no set order. We might have a quantifier always , for example, that implicitly quantifies over all variables in sight . Let us symbolize it by $\forall$ (without attached variables) and take $\forall A$ to indicate that A holds no matter what values are assumed by the free variables occurring in A . The question now arises as to how the algebraist is to understand a sentence such as $\forall(x + y = y + x)$.[4] Clearly, he must understand it in terms of the application of a universality operator to a λ-term constructed from $x + y = y + x$. But which λ-term? There would appear to be only two options: (i) it is a λ-term, $\lambda xy(x + y = y + x)$ or $\lambda yx(x + y = y + x)$, in which the variables attached to the λ-symbol are taken to occur in a set order; (ii) it is a λ-term, $\lambda(x + y = y + x)$, in which the λ-symbol is taken, like the quantifier-symbol $\forall$, to apply to all free variables in sight, though in no set order.

But the first option gives a more specific meaning to the sentence than it actually appears to possess. We must therefore either arbitrarily adopt one interpretation over another or attribute to the sentence an indeterminacy in meaning which it does not seem to have. And one should not think that the relevant order might be given by the standard alphabetic order of the variables or by the order in which they occur in the formula that follows the quantifier, for we might take the variables to be symbols - such as # and * - which are not alphabetized in any given order and

[4]Mathematicians sometimes use the strict equality sign $\equiv$ in this way: $x + y \equiv y + x$, for example, is used to indicate the universal truth of $x + y = y + x$. I might note that a similar difficulty arises for branching quantifiers, since the different branches of the quantifier do not occur in a set order.

the formula might be written in a nonlinear notation from which no set order of occurrence can be discerned.[5]

The second option is clearly more faithful to our semantic intentions. But it plays havoc with the idea of distributing the binders across the logical connectives. For, on pain of reintroducing an arbitrary order on the variables, a λ-term such as $\lambda(x > y)$ will have to symbolize a relation that is neutral between the biased relations symbolized by $\lambda xy(x > y)$ and $\lambda yx(x > y)$.[6] But this means that, when we push the binder λ through $\lambda(x > y \vee y > x)$ in order to obtain the disjunction of the relations symbolized by $\lambda(x > y)$ and $\lambda(y > x)$, we will lose track of the alignment between the variables in the two disjuncts and will therefore be unable to distinguish, in the way we should, between $\lambda(x > y \vee y > x)$ and $\lambda(x > y \vee x > y)$.

A similar difficulty arises from the use of modal and other intensional operators. Suppose that we add an operator $\Box$ for necessity to the symbolism for first-order logic and that we take the interpretation of 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 $\lambda x\Box(x = x)$ and of $\lambda x\Box\exists y(y = x)$. 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 by $\exists y(y = x)$. 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 presumably means, at least under any natural way of understanding the operation of necessitation, that $\lambda x\Box(x = x)$ and of $\lambda x\Box\exists y(y = x)$ will also be modally indistinguishable - which, of course, they are not, since the first is true of any individual whatever while the second is only true of individuals that necessarily exist.

The general 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 a satisfactory account of quantification should be of general import; it should not depend, for its success, on 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 which have no counterpart in our actual understanding of the symbolism. Let me merely give one illustration. When we push the λ-operators through logically complex formulas, we will eventually reach λ-terms of the form $\lambda x_1x_2 \ldots x_mFy_1y_2 \ldots y_n$, in which the binder $\lambda x_1x_2 \ldots x_m$ governs an atomic formula $Fy_1y_2 \ldots y_n$ (and where some of the variables $x_1, x_2, \ldots x_m$ may, of course, be the same as one another and the same as some of the variables $y_1, y_2, \ldots, y_n$). How then are these terms to be interpreted?

---

[5] As with the notation for fractions or determinants: $\frac{x}{y} = \frac{x}{y}$ and $\begin{vmatrix} x & y \\ x & y \end{vmatrix} = \begin{vmatrix} x & y \\ x & y \end{vmatrix}$.

[6] Some accounts of what a neutral relation might be are considered in Fine [00].

Presumably, in keeping with the algebraic approach, we should take $\lambda x_1 x_2 \ldots x_m F y_1 y_2 \ldots y_n$ to signify the result of applying some operation to the relation *F* signified by the predicate F. So, for example, λxFxx will indicate the reflexive version of *F*, while λyxFxy will signify the converse of *F* (at least, if λxyFxy signifies *F* itself). But what of the operation itself? It must presumably be determined on the basis of the relative disposition of the variables $x_1 x_2 \ldots x_m$ in the binder and of the variables $y_1 y_2 \ldots y_n$ in the atomic formula. But there are different ways in which this might be done; and nothing to choose between them. Consider λzxFxxz, for example. We could take this to be the result of first forming the converse λzxyFxyz of λxyzFxyz and then forming the reflexive version λzxFxxz of the converse, or we could take it to be the result of first forming the reflexive version λxzFxxz 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 so, again, we face the awkward choice of making the interpretation either arbitrarily specific or unacceptably indeterminate.

The algebraic approach is best viewed as an attempt to see the symbolism of first-order logic as something which it is not. What it provides, in effect, is a translation from a language with variables to one without variables. A λ-term such as λzxFxxz, for example, may be taken to be equivalent in meaning to a term Refl(Conv(F)), indicating the application of reflexive and converse operations to the relation signified by F; and all appeal to variables in the target language is thereby made to disappear. But in making the transition to a variable-free notation, not only are we forced to make arbitrary decisions about how the translation should go, we thereby loose what is most distinctive about the use of variables. For instead of being treated as devices of reference, albeit of a special sort, they are treated as more or less oblique ways to indicate the application of various operations within a calculus of relations. The problem of understanding our use of variables is not solved, but side-stepped.

§6 The Relational Approach

I now wish to indicate how I think the antimony is to be solved and how a more satisfactory semantics for the symbolism of first-order logic might thereby be developed.

I agree with the autonomous approach in thinking that the formulation of the antimony embodies a false presupposition. But the false presupposition lies not in the attribution of a semantic role to free variables but in the presumption that there is conflict between the sameness in semantic role of x and y, on the one hand, and the difference in semantic role of x, y and x, x, on the other ((SS)′ and (SD)′ above).

There are, I believe, two things that stand in the way of our seeing how these attributions of sameness and difference might be reconciled. The first concerns a possible ambiguity in the notion of semantic role. We have already had occasion to distinguish between sameness or difference of semantic role *across* contexts and *within* a given context. But there is, I believe, another ambiguity in the notion of semantic role that might stand in the way of seeing how reconciliation is to be achieved.

This may be brought to light by considering the following argument against reconciliation. Suppose, in conformity with (SD)′, that the semantic roles of the pair x, y and of the pair x, x are not the same. Then it may be argued that the semantic roles of the individual variables x and y cannot be the same, in contradiction to (SS). For x, when paired with x, has the same semantic role as x, x whereas y, when paired with x, does not have the same semantic role as x, x. The variables therefore differ in respect of whether their pairing with x gives something with the same semantic role as x, x; and this difference, presumably, is a semantic difference.

It is not to be denied that there is a semantic difference of the presumed sort between x and y. But it is not in this sense that we wish to deny that there is a semantic difference between x and y. To see how this is so, let us distinguish between the *intrinsic* (or non-relational) and the *extrinsic* (or relational) semantic features of an expression. The intrinsic semantic features of an expression, in contrast to its extrinsic semantic features, do not concern its semantic relationship to other expressions. Thus it will be an intrinsic semantic feature of the predicate doctor that it is true of doctors but not an intrinsic semantic feature that it is synonymous with physician . Likewise, the intrinsic semantic features of a *pair* of expressions will consist of those semantic relationships between the expressions which do not concern their semantic relationship to yet further expressions. Thus it will be an intrinsic semantic feature of the pair doctor and physician that they are synonymous, though not that they are both synonymous with licensed medical practitioner .

Now what the above argument shows is that if there is an intrinsic semantic difference between the pairs of variables x, y and x, x, then there will be an extrinsic semantic difference between the individual variables x and y themselves (concerning the relationship of each to the variable x). But in asserting that the semantic role of x and y is the same, we only wish to assert that their intrinsic semantic features are the same; and in asserting that the semantic roles of x, y and x, x are different, we mean to assert that their intrinsic semantic features are different. Thus the present difficulty will not arise as long as we always take semantic role to be intrinsic.

The other impediment to achieving reconciliation rests upon a mistake in doctrine rather than upon a failure to recognize a distinction. Let it be granted that the relevant notion of semantic role is both cross-contextual and intrinsic. Still, it might be asked, how can there be a difference in the (intrinsic) semantic relationships holding between each of two pairs of expressions without there being a difference in the intrinsic semantic features of the expressions themselves? Thus given a difference in relationship between the pairs doctor , dentist and doctor , doctor (with the one being synonymous and the other not), there must be a difference in meaning between doctor and dentist . Similarly, given that there is a difference in the semantic relationships holding between the pairs x, y and x, x, there must be a difference in meaning between x and y. Indeed, without such a difference, we would have, in each case, a pair of synonymous expressions; and so it is hard to see how they could give rise to a difference in relationship.

According to this view, there can be no difference in intrinsic semantic relationship without a difference in intrinsic semantic feature. All differences in meaning must be attributable to intrinsic differences; and any attempt to reconcile the attributions of semantic sameness and difference is doomed to failure.

It has to be acknowledged that this view of meaning - what we might call semantical intrinsicalism - is extremely persuasive. But it is false; and a careful examination of the behavior of variables indicates how. Suppose again, to fix our ideas, 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.

What is important to appreciate here is 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 antimony is to be solved. We must allow that any two variables will be semantically the same, even though pairs of identical and of distinct variables are semantically different; and we should, in general, be open to the possibility that the meaning of the expressions of a language is to be given in terms of their semantic relationships to one another.

Formally, the situation is analogous to failures in the identity of indiscernibles. Consider, for example, the distinct but indiscernible spheres of Max Black [70]. Just as there is no intrinsic spatial difference between the two spheres (once we reject absolute space), so there is no intrinsic semantic difference between two variables. And just as there may be an intrinsic spatial difference between two pairs of spheres - since identical spheres will be coincident whereas distinct spheres will not be, so there may be an intrinsic semantical difference between pairs of variables. From this perspective, then, the lesson to be drawn from the antimony is that semantics provides another, though less familiar, example of an aspect of reality in which things can only be distinguished in terms of their relations to one another, and not solely in terms of their intrinsic features.

§7 Relational Semantics for First-order Logic

I now wish to consider how the relational solution to the antimony is capable of yielding a more satisfactory semantics for the symbolism of first-order logic. I deal first with the case of free variables, which is somewhat more straightforward, and then consider the complications that arise from the use of bound variables.

I take it to be the aim of a semantics for a given language to account for the semantic behavior of its expressions. The way 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 is no longer sufficient 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). A semantics should concern itself, then, with the assignment of semantic connections, and not merely with the assignment of semantic values.

Indeed, we should go further still. It is not that semantic enquiry, as it is usually conceived, needs to be *supplemented* by an account of semantic connection. For, given that there are irreducible semantic relationships among expressions, one can reasonably expect that the semantic value of certain complex expressions will depend on the semantic connections among the simpler expressions from which it is derived, and not merely on their semantic values. Thus semantic connection should *replace* semantic value as the principal object of semantic enquiry and the aim of semantics should be to determine the semantic connections among expressions on the basis of the semantic connections among simpler expressions.

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.

If we are to apply this general idea of a relational semantics to first-order logic, we must first have some conception of what we want the semantic connections to be. This may be done, I believe, 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 (obtained under the assignment of 2 to x, 3 to y and 6 to z) and the sequence 6, TRUE, 2 (obtained under the assignment of 4 to x, 2 to y and 2 to z). It should be noted that the semantic connections are entirely non-typographic; they contain no trace of the expressions from which they were derived and the danger of the semantics being implicitly typographic is thereby avoided.

We must now show how to determine the semantic connection on any give 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, in general, the semantic connection on the variables x_1, x_2..., x_n should be taken to be the set of all n-tples a_1, a_2..., a_n of individuals from the domain for which $a_i = a_j$ whenever $x_i = x_j$ $(1 \pm i < j \pm n)$. It is at this point that relationism enters the semantic scene.[7]

The semantics must now be extended to complex terms. 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 nonnegative reals; and the second should have as its 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 ranges of x and x, and similarly for x.y, then we cannot distinguish between them, since the value ranges of x and of y are the same. However, we take the value range of x.x to be a function of 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 comprising all identical pairs of reals and the second comprising 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 give us the result we want.

More generally, let us suppose that we have a complex term $ft_1t_2 \ldots t_n$ and that we wish to determine the semantic connection on a sequence of expressions involving the term. The general

---

[7]The lexical semantics should also let us combine the connection on two sequences of expressions - such x, y and +, 0 - which are independent of one another.

form of such a sequence will be $D_1, D_2 \ldots D_p, ft_1t_2 \ldots t_n, E_1, E_2, \ldots, E_q$ (with p, q ≥ 0). Now, under the standard approach to semantics, the semantic value of the complex term $ft_1t_2 \ldots t_n$ will be determined on the basis of the semantic value f of its function symbol f and the semantic values $a_1, a_2, \ldots a_n$ (under a given assignment) of its argument-terms $t_1, t_2, \ldots, t_n$; and this semantic value will be taken to be the result $f(a_1, a_2, \ldots a_n)$ of applying the function f to the arguments $a_1, a_2, \ldots a_n$. This suggests that the semantic connection on the sequence $D_1, D_2 \ldots D_p, ft_1t_2 \ldots t_n, E_1, E_2, \ldots, E_q$ should be determined on the basis of the semantic connection on $D_1, D_2 \ldots D_p, f, t_1, t_2, \ldots t_n, E_1, E_2, \ldots, E_q$ (with $ft_1t_2 \ldots t_n$ giving way, as before, to $f, t_1, t_2, \ldots, t_n$); and $d_1, d_2, \ldots d_p, b, e_1, e_2, \ldots, e_q$ will belong to the semantic connection on $D_1, D_2 \ldots D_p, ft_1t_2 \ldots t_n, E_1, E_2, \ldots E_q$, i.e. $ft_1t_2 \ldots t_n$ will be capable of taking the value b when $D_1, D_2 \ldots D_p, E_1, E_2, \ldots, E_q$ take the values $d_1, d_2, \ldots d_p, e_1, e_2, \ldots, e_q$, just in case, for some individuals $a_1, a_2, \ldots, a_n$ for which $b = f(a_1, a_2, \ldots a_n)$, $t_1, t_2, \ldots t_n$ are capable of taking the values $a_1, a_2, \ldots a_n$ when $D_1, D_2 \ldots D_p, E_1, E_2, \ldots, E_q$ take the values $d_1, d_2, \ldots d_p, e_1, e_2, \ldots, e_q$, i.e. just in case, for some individuals $a_1, a_2, \ldots a_n$, the n-tple $d_1, d_2, \ldots d_p, f, a_1, a_2, \ldots, a_n, e_1, e_2, \ldots, e_q$ belongs to the semantic connection on $D_1, D_2 \ldots D_p, f, t_1, t_2, \ldots t_n, E_1, E_2, \ldots, E_q$.

The above rule is easily extended to the case of atomic formulas; and a similar rule may be given in the case of truth-functionally complex formulas. The semantic connection on $\neg A$, E, for example, will consist of all those pairs τ', e for which τ, e is a member of the semantic connection on A, E and τ' is the opposite truth-value to τ.

Quantifiers raise some additional problems. The obvious way of evaluating a sequence containing a quantified formula - such as $\exists x A(x)$, E - is in terms of the connection on x, A(x), E. If, for a fixed semantic value of E, A(x) is true for the assignment of some individual from the domain to x , then $\exists x A(x)$ should be taken to be true and otherwise should be taken to be false. Thus the pair TRUE, e will belong to the connection on $\exists x A(x)$, E just in case, for some individual a from the domain, the triple a, TRUE, e belongs to the connection on x, A(x), E (and similarly for the pair FALSE, e).[8]

But consider how such a rule applies to a sequence of the form $\exists x A(x)$, x - say to $\exists x(x > 0)$, x (where the domain of quantification is the set of all natural numbers). The pair TRUE, 0 will belong to the semantic connection on $\exists x(x > 0)$, x just in case, for some natural number n, the triple n, TRUE, 0 belongs to the semantic connection on x, x > 0, x. But since the first and third variables in x, x > 0, x are the same, the first and third components of any triple in its semantic connection will be the same. It follows that no triple of the form n, TRUE, 0 can belong to the semantic connection on x, x > 0, x and so the pair TRUE, 0 will not belong to the semantic connection on $\exists x(x > 0)$, x - contrary to our intentions.

The problem is that we do not want the bound occurrences of the variable x in $\exists x A(x)$, x to be coordinated with the free occurrence. However, our method of evaluation requires that they be coordinated since $\exists x A(x)$, x is evaluated in terms of x, A(x), x. What makes the problem especially acute is that we wish to subject $\exists x(x > 0)$, x to essentially the same method of evaluation as $\exists y(y > 0)$, x, since they are mere notational variants of one another, and yet we also

---

[8] I have assumed that the quantified formula $\exists x A(x)$ contains no free variables. If it does, then we can only properly evaluate the formula by considering its connection to the free variables that it contains; the above rule is in effect relativized to an assignment of individuals to those variables.

want to subject ±y(y > 0), x to the straightforward evaluation in terms of y, y > 0, x.

The way out of the impasse, I believe, is to give up the assumption that all occurrences of the same variable should be treated in the same way. We have so far assumed that different free occurrences of the same variable should always take the same individual as value; and this is, indeed, a reasonable default assumption to make. However, when a free occurrence of a variable was originally bound, only becoming free in the process of evaluation, then we should no longer assume that it is coordinated with those occurrences of the same variable that were either originally free or had their origin in a different quantifier.

This means that, for the purpose of evaluating a sequence of expressions, we should explicitly indicate which of the free occurrences of a given variable are to be coordinated and which not. In the sequence x, x > 0, x, for example, we should distinguish between the cases in which none of the occurrences of x are to be coordinated, in which all are, or in which only two are. (The reader might picture these coordinating links as lines connecting one occurrence of the variable to another.)

This then provides us with the means of evaluating ±x(x > 0), x in terms of x, x > 0, x, in strict analogy with the evaluation of ±y(y > 0), x in terms of y, y > 0, x. However, the variables in the evaluating sequences should be subject to the appropriate pattern of coordination; the first two occurrences of x in x, x > 0, x should be coordinated with one another (though not with the third occurrence) and the first two occurrences of y in y, y > 0, x should likewise be coordinated. In general, whenever we are evaluating a quantified formula ±xA(x) in terms of its components x, A(x) within the context of a sequence, the designated occurrences of x should be taken to be coordinated with one another, though not with any other occurrences of x that may happen to be present in the sequence.

Modest as this proposal might appear to be, its development calls for fundamental revisions in our previous formulation of the semantics. In the first place, the syntactic object of evaluation will not longer be a sequence of expressions but a *coordinated* sequence of expressions. This is a sequence of expressions $E_1, E_2, \ldots E_n$ 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 lexical rule for variables 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. With these two changes in place, the semantics may be stated much as before.

The relational semantics has several clear advantages over its rivals. First and foremost, it embodies a solution to the antimony: 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 on those pairs) are different. The semantics is also more satisfactory, in various ways, *as* a semantics. In contrast to the autonomous approaches, it assigns a semantic role to open expressions; in contrast to the instantial approach, it can be given an extensional formulation; and in contrast to the algebraic approach, it is based upon a credible direct 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 (or some surrogate thereof) into the very identity of the entities that the semantics assigns to the open expressions of the language.

The reader may be worried by the complexity of our account, especially in comparison with the Tarski semantics. But one should bear in mind that the more abstract a semantics for a given language, the more removed it is from the underlying syntax, the more complicated it is likely to be. Indeed, it seems to me that the complexity of our own account is of a desirable sort. For it appears not to be a mere artifact of the approach but to reflect a genuine complexity in our understanding. The relational approach forces us to take seriously the pattern of coordination among the variables; these patterns must be explicitly used as syntactic inputs to the semantic method of evaluation. And it is hard not to believe that their semantic significance, as revealed by the semantics, is not somehow integral to our understanding of the variables. From among the alternatives, it is only the algebraic approach that takes explicit account of the phenomenon of coordination; but it does this in a manner that is completely artificial and not in keeping with the variable s referential role.

§8 Implications

Our approach to the semantics of variables has various broader implications of which I should provide a brief discussion.[9]

It is important, in this connection, to distinguish between the present version of relationism and the more familiar doctrine of semantic holism. Our own version of relationism is meant to serve as a *extension* to the usual forms of representational semantics. We take it to be a legitimate task of semantics to assign referents or senses to expressions and then consider whether semantics, as so conceived, might also require the postulation of semantic relationships. Semantic holism, on the other hand, is meant to serve as a *alternative* for the usual forms of representational semantics. The task of assigning referents or senses to expressions is abandoned in favor of the task of providing an account of their inferential or conceptual role. Thus the two approaches are relational in very different ways, with the one staying squarely within the representational framework and with the other locating the relational ideas outside of that framework.

As I have already indicated, the proper semantic treatment of variables provides a clear and convincing refutation of the doctrine of semantic intrinsicalism. But semantic relationism might not be so interesting, even as a general thesis, if the case of variables afforded the only instance of its truth. There are, however, many other ways in which the doctrine can be applied.

One of the most interesting, to my mind, concerns the distinction between what what one might call strict and accidental coreference. A Millian will claim that two coreferential proper names, such as Cicero and Tully , will have the same intrinsic semantic features, since there is no more to their meaning than their referents. However, it is still possible, I believe, for the Millian to distinguish the semantic relationship between Cicero and Cicero (which is one of *strict* coreference) from the semantic relationship between Cicero and Tully (which is one of *accidental* coreference). Thus strict (or accidental) coreference will be semantic relationships for the Millian which are not grounded in the intrinsic semantic features of the expressions

[9] I hope to discuss these more extensively in a forthcoming book Reference, Relation and Meaning to be published by Basil Blackwell.

between which they hold. Moreover, once this distinction is made, the Millian has new means of dealing with Frege s puzzle ([52]), with Kripke s puzzle concerning belief ([79]), and with other general issues concerning the relationship between semantics and cognition.

A related distinction may be drawn between *strict* and *accidental* synonymy. Thus doctor and doctor will be strictly synonymous, while doctor and physician will only be accidentally synonymous. This distinction should be accepted by the Fregean., I believe, and not merely the Millian; and it provides us with a new, and more plausible, way of dealing with Mates puzzle ([72]) and with Moore s paradox of analysis.

Relational ideas are also important to the project of constructing a detailed semantics for ordinary language. The most obvious application is to anaphora, which should be treated as a form of strict coreference (or synonymy). But there are numerous correlative terms and constructions which also appear to call for a relational treatment. Coordinated uses of Mr Smith and Mrs Smith , for example, should be taken to refer to husband and wife; and in a construction such as the more the merrier , it is clear that the semantic value of the more should somehow be coordinated with the semantic value of the merrier .

Finally, there are further applications to be made to logic itself. For the apparatus of coordination provides us with much greater flexibility in formulating different forms of semantics for the use of variables. It can accommodate a form of dynamic semantics, for example, one in which it is possible to quantify out; and it enables us to provide a much more plausible account of referential quantification into intensional contexts than the standard approaches, since coordination between variables can now be made to track significant semantic distinctions. Thus we see that, quite apart from its general relevance to the philosophy of language, semantic relationism has detailed implications for the development of semantics in a wide variety of areas.[10]

References

Bealer G., [79] Theories of Properties, Relations, and Propositions , Journal of Philosophy 76, 634-648.

Bealer G., [82] Quality and Concept , Oxford: Clarendon Press.

Bealer G., [83] Completeness in the Theory of Properties, Relations and Propositions , Journal of Symbolic Logic, vol. 48, no. 2, 415-426.

Black M., [70] The Identity of Indiscernibles , in Universals and Particulars (ed. M. Loux), Garden City, NY: Doubleday, 204-16.

Dummett M., [73] Frege: Philosophy of Language , London: Duckworth.

Fine K., [84] A Defence of Arbitrary Objects , Aristotelian Society, Supplementary Volume 57 (1983), 55-77, Reprinted in Varieties of Formal Semantics , (ed. F. Landman and F.

[10] I should like to thank the members of seminars at St Andrews University, New York University and UCLA for helpful discussion. I am especially grateful to Ruth Chang and Louis Derosset for their comments on the written version of the paper.

Veldman), Dordrecht-Cinnaminson: Foris Publications (1984).
Fine K., [85] Reasoning with Arbitrary Objects , Oxford: Blackwell.
Fine K., [89] The Problem of De Re Modality , in Almog et al [89], 197- 272.
Fine K., [98] A Defense of Cantorian Abstraction , Journal of Philosophy, 98/12, 599-634.
Fine K., [00] Neutral Relations , Philosophical Review, vol. 109, no. 1, 1-33.
Frege G., [52] On Sense and Reference , translation of Uber Sinn und Bedeutung , in Translations from the Philosophical Writings of Gottlob Frege (eds. P. T. Geach and M. Black), Oxford: Clarendon Press.
Kripke S., [79] A Puzzle about Belief , in Meaning and Use (ed. A. Margalit), Dordrecht: Reidel.
Mates B., [72] Synonymity in Semantics and the Philosophy of Language (ed. L. Linsky), Urbana: Univ. of Illinois Press, 111-38.
Quine W. V. O. [52] Methods of Logic , London: Routledge & Kegan Paul.
Tarski A., [36] Der Wahrheitsbegriff in den Formalisierten Sprachen , Studia Philosophica I, 261-405, originally published in 1933, reprinted in Logic, Semantics and Metamathematics , Oxford, 1955.
Wittgenstein L., [1922] Tractatus-Logico-Philosophicus , London: Routledge-Kegan.