II. The Proper Interpretation of Variables

We have a fundamental representational device - indeed, *the* fundamental device for conveying statements of generality, and yet when we attempt to provide a systematic account of its semantic function we appear to encounter insuperable difficulties. What makes the problem so exasperating as well as baffling is that we feel the solution should be within easy reach. There is, after all, something completely transparent about our understanding of variables, their role is to represent a given range of values; and so it should be evident from our understanding how a semantics governing their use should proceed. There should be no need either to invent a new kind of role for the variables or to give up on generally recognized principles for the construction of a semantics.

I would like to suggest in this chapter that the fundamental source of our difficulties is the assumption that semantics should take the form of assigning semantic values. We should recognize instead that semantics may take the form of assigning semantic *relationships* to expressions; it is these relationships, rather than the values, that should be the principal object of semantic enquiry. Once we make this shift, then I think that it will be found that all of the difficulties we previously encountered may be resolved.

I begin by presenting a particularly simple version of our problem, which I call the 'antimony of the variable', and am led to conclude, on the basis of an analysis of this antimony, that the doctrine of semantic relationism might provide us with a solution (§1). I then argue that the doctrine of semantic relationism *must* be correct in application to variables and that therefore any standard value-oriented version of semantics for the variables will be mistaken (§2). However, this still leaves us with the task of developing an adequate relational account. I first present a relational semantics for free variables (§3) but, unfortunately, the extension of the semantics to bound variables is not altogether straightforward. The most natural proposals fall prey to a subtle form of typographic bias and, after criticizing these proposals (§4), I show how a more flexible relational approach is able to avoid the bias (§5). The resulting view has significant implications, both for our understanding of predicate logic and for our general conception of the kind of framework to which its syntax and semantics should be taken to belong

(§6).

§1. The Doctrine of Semantic Relationism

We have attempted to find a satisfactory semantics for our use of variables, one which conforms to the general requirements of compositionality and determination and which is faithful to our intuitive understanding of their role. But all of the obvious suggestions have come to grief in one way or another. The first two - the contextual and the dismissive - fail properly to respect one or other of the requirements. The other two - the instantial and the algebraic - perhaps do better at satisfying the general requirements, but the former is non-extensional while the latter is unbelievable, with neither providing a fully adequate account of the semantics for open expressions.

I believe that the problem runs very deep and cannot even be solved in the terms in which it has been stated. So let us reconsider the problem in the hope of understanding how it might be reconceived. At the heart of all our difficulties is an apparent conflict in the semantics of variables. On the one hand, given two variables 'x' and 'y' that range over the same domain of objects, we want to say that they are semantically the same, since 'x > 0', let us say, is merely a notational variant of 'y > 0'. On the other hand, we also want to say that the variables are semantically different, since it is clear that they play distinct semantical roles in 'x > y'. In other words, we wish to say:

(SI) Any two variables (ranging over a given domain of objects) are semantically the same; and

(SD) Any two variables (ranging over a given domain of objects) are semantically different.

Yet clearly we cannot univocally say both, given that there are at least two variables!

We might call this 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.' The solutions we have considered tend to opt for one or other side in this conflict and it is, I believe, only by seeing how both sides might be reconciled that we will be in a proper position to formulate a more adequate semantics.

Clearly, any resolution of the conflict must begin by locating an ambiguity in the phrase 'semantically the same'. Whatever the sense is in which the variables x and y are semantically the same in 'x > 0' and 'y > 0' is not 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 the notion of context.[1] The variable 'x' in the context of the formula 'x > 0'' is semantically the same as 'y' in the context of another formula 'y > 0' (though we might well say that there were semantically distinguishable if both formulas occurred within the context of a single proof). On the other hand, within the context of a single formula, such as 'x > y', the two variables *will* be semantically different. Thus (SI) will hold for semantic sameness in the cross-contextual sense and (SD) for semantic sameness in the intra-contextual sense.

But our problems are not yet over, since it is clear that the two notions of semantic sameness cannot be completely divorced from one another. Why do we say that x and y are semantically different 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. In other words, the *internal* (or intra-contextual) semantic difference of the variables within the formula 'x > y' can be represented as the *external* (or cross-contextual) semantic difference of the pair of variables x, y in 'x > y' from the pair of variables x, x in 'x >x'. In general, to say that x and y are internally different, in the intended sense, is just to say that the pair of variables x, y is externally different from the pair of variables x, x. We may therefore state (SI) as before and state (SD) in the form:

(SD)′ the pair of variables x, y is semantically different from the pair of variables x, x, using an *univocal* external notion of semantic sameness throughout.

There is still no explicit contradiction. But there is a difficulty in seeing how (SI) and (SD)′ might both be true. For how can there be an external semantic difference between the pair of variables x, y and the pair x, x unless there is an external semantic difference between the variables y and x themselves? After all, if there is an external semantic difference between x, y

[1]Though not, of course, in the sense in which assignments might be thought to constitute a context.

and x, x, then there must be a semantic difference between x and y. And so why should this difference only be manifest when the variables are used in the same context and not when they are used in different contexts? Why is it not just a general context-insensitive difference?

We have here an especially pure and acute form of our problem. For it is stated entirely in terms of variables, without the aid of any other expressions of the language. And we may effectively dispense with any appeal to Compositionality or Determination, since the notion of variables as notational variants may be used to sustain (SI) and the notion of distinct variables as ‘independent’ loci of values may be used to sustain (SD)′. The antimony of the variable is staring us straight in the face.

Before attempting to show how this further conflict might be resolved, we need to enquire more deeply into why it should be thought that the assumptions (SI) and (SD)′ are in conflict. The reason, I believe, is that we implicitly take for granted a certain principle connecting the semantic relationships among expressions to their semantic features. According to this principle, which we dub ‘Intrinsicalism’:

(Intr) Any semantical relationship among expressions must be grounded in their semantical features.

In stating the antimony, it was presupposed that any external distinction between two expressions must be made in terms of the semantic *features* of the two expressions and that any internal distinction between two expressions E and F - or, equivalently, any external distinction between E, F and E, E - must be made in terms of a semantic *relationship*, one that holds between E and F but not between E and E. Suppose now that there is an internal semantic difference between x and y - or, equivalently, an external semantic difference between x, y and x, x. There will then be a semantic relation, call it *S*, that holds between x and y but not between x and x. By (Intr), this semantic relationship must be grounded in semantic features of the variables. There must, in other words, be a semantic feature *P* of x and a semantic feature *Q* of y whose presence in any two variables will guarantee that they are related by *S*. But then x cannot possess the semantic feature *Q* possessed by y, since otherwise the semantic relation *S* would hold between x and x. And so given that y possesses *Q* and x does not, it follows that there is an external semantic difference between x and y.[2] The conflict between (SI) and (SD)′ is thereby sustained.

It is important to avoid two extreme reactions towards (Intr), both of which serve to trivialize what is at issue. According to the first, it might be argued that the principle - or, at least, the application we wish to make of it - cannot sensibly be denied. For suppose that the pair of variables y, x stand in a semantical relationship to which the pair of variables x, x do not stand. It will then be trivial that there is a semantic feature possessed by y but not by x, since y will possess the feature of standing in the given semantic relationship to x, whereas x will not.

It must, of course, be granted that there is this difference between x and y. However, this difference is not, in the relevant sense, a difference in *semantic feature*. For we require that the semantic features of an expression should be intrinsic to the expression and not concern its relationship to other expressions (and similarly, we require that a semantic relationship between two expressions E and F should be intrinsic to the pair of them and not concern their relationship to other expressions). That the features are intrinsic in this sense might be taken to be implicit, of course, in the very notion of a *feature*. But it might also quite plausibly be taken to be implicit in the requirement that the features should be *semantic*. For if they are semantic (or perhaps we should say *purely* semantic), then they will not involve the kind of typographic bias that is invoked in saying how one expression relates to another. But regardless of whether or not this is so, we shall insist that the semantic features spoken of in (Intr) should be intrinsic to their bearers; and there will be then be no way in which the principle might be shown to be trivially true.

The other extreme reaction is to maintain, not that (Intr) is trivially true, but that it is trivially - or familiarly - false. For have not the holists been telling us for a long time that the semantics for a language must be understood by reference to the semantical relationships among its different expressions, and not merely by reference to the individual semantical features of the

[2] It is apparent from this ‘proof’ that we have stated (Intr) in an excessively strong form. All that we strictly require is that when *S* is a semantic relation between expressions E and F there should be semantic features *P* and *Q* of E and F respectively which are such that any expressions possessing the respective features *P* and *Q* will be related by *S*.

expressions themselves. And so does not the falsity of (Intr) simply follow from the truth of holism?

I believe, however, that this is to seriously misrepresent what is here at issue. The underlying dispute between the holists and their opponents is over what constitutes a proper account of the semantics of a given language. 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 a representational semantics, were it to be given, would be atomistic; and it is only because the holists wish to adopt an alternative form of semantics that they take their position to be even remotely plausible.

It is on this point of agreement, however, that we wish to locate the current dispute. For our concern is to reject (Intr), not for an inferential form of semantics, or even for semantics in general, but for a representational form of semantics. Thus our disagreement is as much with the holists as with their opponents.

We see that, under a proper understanding of (Intr), the semantics should be taken to be representational rather than inferential in form and the semantic features should be taken to be intrinsic rather than extrinsic to their expressions. However, once (Intr) is understood in this way, it seems hard to deny. For if our account of the semantics for a given language is in terms of its representational role, then what room is there for any kind of semantic relationship that is not grounded in the representational features of the expressions? How can representational function of language be distributed, as it were, across different expressions?

We can bring out how plausible (Intr) is under the present construal by means of an analogy with the substantivalist conception of space. The substantivalist believes that the proper account of the spatial facts is to given by assigning a location to each spatial object; the relationist, on the other hand, thinks that the proper account is to be given by specifying the spatial relations among the objects themselves. Now just as the substantivalist believes that the spatial facts are to be specified by assigning positions to spatial objects, so the representationalist thinks that the semantic facts are to be specified by assigning ‘meanings’ or semantic values to expressions; and just as the relationist thinks that the spatial facts are to be specified by describing the spatial relationships among different spatial objects, so the inferentialist thinks that the semantic facts are to be specified by stating the semantic or ‘inferential’ relationships among different expressions. But imagine now a substantivalist who thinks that the proper account of the spatial facts requires, not only an assignment of location to spatial objects, but also the specification of certain spatial relationships among those objects themselves. This would appear to be metaphysically weird for surely, one might think, the assignment of location should be sufficient in itself to fix the spatial facts (along, of course, with the geometry that is intrinsic to the locations). But it is exactly this weird hybrid position that is analogous to the semantical position that it is here under consideration. For it is being suggested that the representational facts are not simply given by the meanings or semantic values of different expressions (and the relationships among those meanings or semantic values) but also by the semantic relationships among the expressions themselves.

In order to distinguish the present position from the holist’s, let us call it ‘Semantic Relationism’; and let us call the opposing view, ‘Semantic Intrinsicalism’. Despite its very great plausibility, I wish to argue that Intrinsicalism is false. Indeed, it seems to me that the case of variables provides a convincing demonstration that it is false; and it is only once we have rejected the intrinsicalist’s view that we will be in a position to solve the antimony of the variable and to develop a satisfactory semantics for predicate logic.

§2. The Defense of Semantic Relationism

We may show that Semantic Intrinsicalism is false for variables simply at the level of the lexicon, i.e. by considering the semantics for variables without any regard for how they might be used in more complex expressions.

Suppose that our language contains the variables v_1, v_2, v_3, ... How then, in the semantics for the language, should we describe their representational role? What, in a word, is their lexical semantics?

Now we should certainly specify the values each variable can assume. For simplicity, let us suppose that each variable draws its values from a fixed domain of objects *D* (though in a many-sorted language, this need not be so).[3] The following rule should then be part of the lexical semantics:

Range for each variable x, the object *a* is a value of x iff it belongs to the domain *D*.

One might now think that one was done; and, indeed, many authors in describing the intended interpretation of the variables do no more than specify their range. However, more does need to be said. For we need to specify not only which values each variable can assume when considered on its own but also which values several variables can assume when considered together. We need to specify, for example, not only that v_1 can take Plato as a value, say, and that v_2 can take Aristotle as a value but also that v_1 and v_2 can collectively take Plato and Aristotle as values.

The need to go beyond a simple specification of the range of values has two significant consequences for the formulation of the semantics. The first is that the primitive semantic predicate by means of which we specify the representational role of variables cannot just be taken to be the two-place predicate that is true of a variable and each of its values. Instead, we should countenance, for each natural number n > 0, a 2n-place predicate that is true of n variables $x_1, x_2, ..., x_n$ and n objects $a_1, a_2, ..., a_n$ when the variables $x_1, x_2, ..., x_n$ can collectively assume the objects $a_1, a_2, ..., a_n$ as values (the previous semantic primitive being the special case in which n = 1).[4] The second consequence is that we need to specify what the application of the new semantical primitive should be. In place of Range above, we should now have the following rule:

Coordination for any variables $x_1, x_2, ..., x_n$ and objects $a_1, a_2, ..., a_n$, $x_1, x_2, ..., x_n$ can take the objects $a_1, a_2, ..., a_n$ as values iff (i) each of $a_1, a_2, ..., a_n$ is in the domain *D* and (ii) the objects a_i and a_j are the same whenever the variables x_i and x_j are the same, $1 \leq i < j \leq n$.

(i) corresponds to our previous domain condition; and (ii) is a condition governing the

---

[3] Of course, there is the question as to whether we think of the domain as given by a class or by something more intensional. Such questions are of no importance here.

[4] We might, of course, appeal to a single predicate of multiple adicity in place of a multiplicity of several predicates of fixed adicity.

independence in value of distinct variables and the coordination in value of identical variables. Together, they are plausibly taken to provide a complete lexical semantics for the variables of the language; nothing else, at the level of the lexicon, need be said concerning their representational role.[5]

We are now in a position to provide some kind of 'proof' that Semantic Intrinsicalism is false. For two distinct variables x and y will be (*value*) *independent* in the sense that for any objects *a* and *b* of the domain it is possible, simultaneously, for x to assume the value *a* and for y to assume the value *b*. A single variable x, on the other hand, will not be independent with itself for, assuming that there are least two objects *a* and *b* in the domain, it will not be possible, simultaneously, for x to assume the value *a* and for x to assume the value *b*. If now we attempt to describe the intrinsic semantic features of two variables, they will be the same; all that we can in effect say about any one of them is that it takes the values that it does. It therefore follows that the semantic relation of independence between two variables x and y is not grounded in their intrinsic semantic features; for, if it were, the relation would also hold between x and x.

In the above 'proof', I have simply taken for granted that independence is an intrinsic semantic relation between two variables and that any two variables will have the same intrinsic semantic features. This is very intuitive but, in the present context, it is possible to state a perfectly precise criterion for what it is for a semantic relation or feature to be intrinsic. For this simply means that it is 'purely qualitative'; its application does not turn on the identity of any variables (beyond those, of course, to which it is applied). It is then evident, under this criterion, that *independence* is an intrinsic relation and that any two variables will have the same intrinsic features.[6]

---

[5]Assuming, of course, that we only assign values to finitely many variables at any one time. The framework must be appropriately extended if we allow the simultaneous assignment of values to an infinitude of variables. We shall later need to see the need to qualify this remark, but the qualification will not affect what is here at issue.

[6]Formally, an intrinsic feature or relation is one that is invariant under the permutations of the variables. The failure of Intrinsicalism then turns on the fact that every permutation of variables induces an automorphism on the corresponding semantic structure. We should note that independence of m variables does not imply independence of n variables, n > m, and so we

The situation described above is formally analogous to Black's case of the two indiscernible spheres (Black [70]). We might think of Coordination as describing a miniature semantic universe, whose sole objects are the variables and the values and whose sole relations are those that hold between a given number of variables and their values. Within such a world, the variables will be intrinsically indiscernible, just as Black's spheres are intrinsically indiscernible. There will, however, be relational differences between the variables just as there may be relational differences between the spheres. Thus one variable may be 'independent' of another, even though that other variable is not independent of itself - just as one sphere may be spatially disjoint from the other, even though that other sphere is not spatially disjoint from itself. From the perspective of general metaphysics, the failure of intrinsicalism within the semantic sphere is no more surprising that its failure in other, more familiar contexts.

Now it might be thought that the proposed refutation of Intrinsicalism is somehow suspect, for why should we think that the 'independence' of distinct variables is genuinely a separate fact over and above their taking the values that they do? Surely it should follow from principles constitutive of what it is to be an assignment that distinct variables will assume values independently from one another; and so it should not be necessary to state this as a specific semantic feature of the variables themselves. Now even if there were principle that guaranteed the independence of distinct variables, I suspect that issues of dependence and independence would re-emerge at the level of complex expressions (such as Px and Py) that could not be handled in a similar way. But it is worth noting that, though there may be some principles that are genuinely constitutive of the notion of assignment[7], there are none that can plausibly be taken to govern their independence; for there is nothing in the notion of what it is to assign values to variables which requires that the values should be capable of being assigned to the variables

---

will need to posit irreducibly semantic relationships of arbitrarily high adicity.

[7] Two such principles are:

(i) Truncation: if $x_1, x_2, ..., x_n$ can take $a_1, a_2, ..., a_n$ as values, then $x_2, ..., x_n$ can take $a_2, ..., a_n$ as values; and

(ii) Permutation: if $x_1, x_2, ..., x_n$ can take $a_1, a_2, ..., a_n$ as values, then $x_{f(1)}, x_{f(2)}, ..., x_{f(n)}$ can take $a_{f(1)}, a_{f(2)}, ..., a_{f(n)}$ as values, where f is any permutation of $\{1, 2, ..., n\}$.

independently of one another.

Indeed, there are many uses of variables, both hypothetical and actual, that are in contravention of such a principle. Thus Wittgenstein ([1922]) has proposed that distinct variables should always assume distinct values; and this is not an altogether uncommon usage. (I have seen a state mathematics exam, for example, in which it was implicitly taken for granted that {x, y} is a two-element set). Or again, in notating a game of chess, we will use the same symbol for a pawn in different squares to stand for *different* pawns, and not the same pawn. Or again, there are what one might call 'correlative' uses of variables. Where 'A' is a variable for a formula, for example, we might use '**A**' as a variable for its Goedel number. In all of these cases, the possibilities for the independent assumption of values are curtailed in one way or another; and so the fact that the variables in some symbolism *are* allowed to receive values independently of one another should be taken to be a separate and, indeed, quite distinctive fact governing their use.

Once Intrinsicalism is given up, the antimony of the variable is solved - for we no longer have any conflict between the external semantic sameness and the internal semantic difference between two variables. Indeed, we see clearly, once we think specifically in terms of the semantic behavior of variables and not of general semantic doctrine, how two variables can be both semantically the same and semantically different, since the intrinsic semantic features of any two variables will be the same even though there are semantic relations of independence that will hold between them though not between any identical pair of variables. It was only the general adherence to false though plausible semantic doctrine that prevented us from seeing - or, at least, from properly describing - the specific semantic facts concerning variables for what they are.

§3. Relational Semantics for Free Variables

We have felt obliged to give up the doctrine of Semantic Intrinsicalism and I now wish to consider what consequences this has both for our general conception of semantics and for the semantics of predicate logic in particular. These consequences are quite considerable - the general nature of semantics and the semantics for predicate logic must both be fundamentally reconceived; and we shall therefore find it convenient to consider the relatively straightforward case in which the language contains only free variables before turning to the complications that arise from allowing the use of bound variables.

An immediate casualty of abandoning Intrinsicalism is Compositionality. According to this requirement, at least as we have stated it, when the complex expression E is (syntactically) derived from the simpler expressions $E_{1,} E_2, ..., E_n$, the semantic value *e* of E should be (semantically) derivable from the respective semantic values $e_1, e_2, ..., e_n$ of $E_{1,} E_2, ..., E_n$. However, such a formulation is inappropriate once we abandon the doctrine of Semantical Intrinsicalism. For we must then allow that there may be semantic relationships among the expressions $E_{1,} E_2, ..., E_n$, which are not grounded in the semantic values of those very expressions. So when E is syntactically derived from $E_{1,} E_2, ..., E_n$,, we should allow that the semantic value of might depend, not merely upon the semantic values of $E_{1,} E_2, ..., E_n$,, but also upon the semantic relationships among them. Thus the semantic value of an expression should be taken to depend, in general, upon the semantic relationships among the expressions from which it is derived (including, as a special case, the semantic values of the expressions themselves).

But even this does not go far enough. For a semantics should be inductive; and so the semantic relationships among expressions must also somehow be determined. The simplest, and most natural, way to do this is through the syntactic analysis of the component expressions. Thus when E is syntactically derived from $F_{1,} F_2, ..., F_m$, the semantic relationships among E, $E_{1,} E_2, ...,$ E_n should be determinable on the basis of the semantic relationships among $F_{1,} F_2, ..., F_{m,}$ $E_1, E_2, ..., E_n$ (E in the sequence gives way to $F_{1,} F_2, ..., F_m$).[8] Thus the semantic relationships among expressions (and not just their semantic values) should now be taken to depend, in general, upon the semantic relationships among simpler expressions.

There may be many different semantic relationships that hold among some given expressions and, in order to achieve a tidy formulation, it will be convenient to wrap these up

---

[8] Note that the resulting syntactic analysis will almost inevitably be ambiguous for, given a sequence $E_{1,} E_2, ..., E_n$ containing several complex expressions, we shall have a choice as to which should first be subject to analysis. I am inclined to think that this ambiguity is real and corresponds, roughly speaking, to the 'order' in which an expression might be evaluated. However, it is clear that we should in general treat the different semantic evaluations that ride on these different syntactic analyses as being, in some sense, strongly equivalent.

into a single entity, which we will call their *semantic connection*. Thus the semantic connection between expressions is, in effect, the sum-total of the semantic relationships that hold among them (including, of course, their semantic values). The semantic connection among expressions could in principle be very complex, but I think that it will found in practice that there is often a simple and unified way of conceiving what it should be.

The aim of semantics should now be taken to be the determination of the semantic connections among expressions on the basis of the semantic connections among simpler expressions. Thus *semantic connection* should replace *semantic value* as the principal object of semantic enquiry; and just as the semantic evaluation of a complex expression is naturally taken to bottom out in semantic values of the lexical items under the standard form of semantics, so it is naturally taken to bottom out in the semantical connections among lexical items under the present form of semantics.[9] We are thereby led to reconceive the general nature of semantic enquiry - and by consideration of what one might have thought was one of the most well understood forms of symbolism!

Although Compositionality - as we have formulated it - must now be abandoned, there is still a general sense in which it may be taken to hold. For the semantic facts concerning certain expressions can still be taken to derive from the semantic facts concerning simpler expressions and, in particular, the 'meaning' of an expression can still be taken to derive from 'the meaning of its constituents', as long as this phrase is understood in a collective sense to embrace the meaning relationships among the different constituents and not just their individual meanings. Our previous formulation of Compositionality can now be seen to be the product of this more general formulation and the intrinsicalist doctrine that the collective meaning of different expressions should be grounded in their individual meanings. Thus Compositionality may still be

---

[9] One could, of course, treat a sequence of expressions of the given language as itself an expression of the language; and the semantic connection among the component expressions could then be identified with the semantic value of the sequential expression that they compose. In this way, one could reduce the connection-based semantics to the value-based semantics. But such an identification would be highly artificial, especially in regard to the items in the lexicon. For any sequence of lexical items upon which one wished to confer a semantic connection would itself have to be treated as a lexical item.

accepted once its formulation is divorced from intrinsicalist presuppositions.

What of the other desiderata - Correlation, Determination and Neutrality? We see that our previous notion of 'semantic difference' should now be taken to be ambiguous, since it can mean either *external* or *internal* difference (the distinction did not arise under the intrinsicalist point of view). If it means external difference, then Correlation and Determination can still plausibly be taken to hold. If on the other hand, it means internal difference, then these principles should be given up, since we cannot expect in general to account for the internal or 'relational' difference between two expressions in terms of their semantic values, i.e. their intrinsic semantic features. However, this is not to abandon the general explanatory aims of semantics, for we should still expect the semantic connection between expressions to account for the presence of absence of an internal semantic difference. Neutrality, on the other hand, should hold in unqualified form. Indeed, not only can we reasonably demand neutrality of the semantic values, we may reasonably demand it of the semantic connections as well. Thus we see that all of the desiderata should still be endorsed once they have been modified to accommodate the substitution of connections for values and the possible ambiguity in the notion of semantic difference.

With this general discussion concluded, let us now consider the implications of relationism for the semantics of predicate logic (absent all devices of quantification). We may follow Tarski in adopting the standard syntax for the symbolism of predicate logic. In particular, each complex term $ft_1t_2 \ldots t_n$ will be derived from its function symbol f and argument-terms t_1, t_2, ..., t_n and each atomic formula $Pt_1t_2 \ldots t_m$ will be derived from its predicate symbol P and subject-terms $t_1, t_2, \ldots , t_n$. Thus there is no need to go against the natural syntactic grain of the symbolism of predicate logic, as in the instantial and algebraic approaches.

We shall find it helpful to distinguish between two types of semantic value for expressions, which I shall the *concrete value* and the *value range*. The concrete value of a closed term will be its denotation; for a closed formula, it will be its truth-value, $\top$ or $\bot$; for an n-place predicate symbol, it will be its n-ary extension over the domain *D;* and for an n-place function symbol, it will be the corresponding n-ary function from *D* into *D*. We may also talk, by obvious extension, of the concrete value of an expression, open or closed, relative to an assignment of objects to its free variables. For example, the concrete value of 2.x will be 6 under the assignment of the number 3 to x.

The notion of a value range is less familiar. This applies to all expressions, open or closed, without reference to an underlying assignment. The semantic value of an expression in this sense will be a *set* of concrete values - specifically, the set of all those concrete values that the expression is capable of assuming under the different assignments of objects to the variables. Thus the value range of x will be the set of all objects in the domain; of 2.x it will be the set of all even numbers (where x ranges over the nonnegative integers); and of x > 0 it will be set of truth-values {$\top$, $\bot$}, since the formula is true when x is positive and false when x is zero.

The idea behind this conception of value is that the semantic role of an open expression is simply to represent a range of values. Thus x represents all objects from the domain, 2x represents all even numbers and, similarly though somewhat less naturally, the formula x > 0 represents the truth-values $\top$ and $\bot$.[10] The value range, as we have defined it, therefore contains far less information than either the Tarski- or the Frege-value, since nothing is said about the connection between the 'output' values, which belong to the value range, and the 'input' values, which are assigned to the variables.

The semantic values of individual expressions, in either sense, may be extended to sequences of expressions. Suppose, first, that $E_{1,}$ E_2, ..., E_n is a sequence of *closed* expressions (each could be a closed term, a closed formula, a predicate symbol, or a function symbol). Then the concrete semantic value of the sequence will be the n-tple of entities ($e_{1,}$ e_2, ..., e_n), where $e_{1,}$ e_2, ..., e_n are the concrete semantic values of the respective expressions $E_{1,}$ E_2, ..., E_n. For example, the concrete value of 2+3, 2 > 3 will be the pair of entities (5, $\bot$).   As before, we may also talk by obvious extension of the concrete value of a sequence of expressions, open or closed, relative to an assignment of objects to its free values. For example, the concrete value of 2.x, x > 0 under the assignment of the number 3 to x will be the pair (6, $\top$).

Suppose now that $E_{1,}$ E_2, ..., E_n is a sequence of expressions, open or closed. Then the

[10]The open expression x > 0 can naturally be taken to represent all propositions to the effect that a given number exceeds 0. It is only when we 'extensionalize' these values that we end up with the current somewhat artificial conception of the expressions's semantic role.

value range, or *semantic connection*, on $E_{1,} E_2, ..., E_n$ will be a set of concrete values on $E_{1,} E_2, ...,$ E_n - specifically, the set of concrete values that the sequence of expressions is capable of assuming under the different assignments of objects to its free variables. For example, the semantic connection on the sequence x, 2x will be the set of ordered pairs of the form (n, 2n).

Just as we may take an open expression to represent a range of values, so we may take a sequence of expressions to represent an interdependent range of values, with the values taken by some of the expressions constraining the values taken by the others. The semantic connection on a sequence of expressions then tells us how their individual values are connected. We may call the members of a semantic connection on the sequence$E_{1,} E_2, ..., E_n$, the sequence's (*admissible*) *assignments*. In the special case in which the expressions $E_{1,} E_2, ..., E_n$ are the distinct variables $x_{1,} x_2, ..., x_n$, the admissible assignments will just be assignments in the sense of Tarski. Thus the notion of a semantic connection, or of an admissible assignment, can simply be regarded as a generalization of Tarski's notion of an assignment.

The fundamental task of a semantics for predicate logic, under the relational approach, is to determine the semantic connections on its expressions. As with a standard value-based semantics, there will be two kinds of rule: those for determining the semantic connections among items in the lexicon; and those for determining the semantic connections among more complex expressions (or, to be more accurate, among sequences of expressions at least one of which is complex). Recall that the lexicon of our language will consist of predicate symbols, function symbols, and variables.[11] The semantical rules for the lexicon will assign an appropriate kind of extension to each predicate symbol and an appropriate kind of function to each function symbol, as is standard. But they will not simply assign a semantic value to each variable; they must also assign a semantic connection to an arbitrary sequence of variables. The coordination principle, mentioned before, tells us how this is to be done:

Coordination Rule $[x_1, x_2, ..., x_n] = \{(a_1, a_2, ..., a_n): a_1, a_2, ..., a_n \in D$ and $a_i = a_j$ whenever $x_i = x_j, 1 \leq i < j \leq n\}$.

Note that the sequence of variables $x_1, x_2, ..., x_n$ may contain repetitions and that the coordination

---

[11] We agreed to incorporate the logical constants into the syntax, but there would be no difficulty in also making them part of the lexicon.

rule then requires that the same object be assigned to repeated occurrences of the same variable.

It is at this critical point that relationism enters the semantic scene; for we require the specification of a semantic connection, and not merely a semantic value. The previous semantical proposals can all be regarded as more or less devious and unacceptable ways of attempting to do without a semantic connection among the variables. The instantial and algebraic approaches do not even admit variables into the lexicon. The Tarski approach does, but assigns them a typographically biased value. Indeed, this can be seen in retrospect to be inevitable; for if we attempt to represent the semantical relations among the variables as semantic features, then we will be forced to make those features typographic.

The semantical rules for the lexicon are not yet complete. For we should specify what semantic connections are to hold on *any* sequence of lexical items - whether they be predicate symbol, function symbols, variables or some combination of these items. So suppose that $E_1, E_2, \ldots, E_n$ is a sequence of variables, predicate symbols and function symbos (in any order). Then $(e_1, e_2, \ldots, e_n)$ will be an admissible assignment for $E_1, E_2, \ldots, E_n$ just in case two conditions are met: first, e_i should be an admissible assignment for E_i whenever E_i is a predicate or function symbol; and second, if $E_{k1}, E_{k2}, \ldots, E_{km}$ is the subsequence of variables in $E_1, E_2, \ldots, E_n$, then the corresponding subsequence $(e_{k1}, e_{k2}, \ldots, e_{km})$ of $(e_1, e_2, \ldots, e_n)$ should be an admissible assignment for $E_{k1}, E_{k2}, \ldots, E_{km}$. What this means, in effect, is that we take there to be no nontrivial semantic connections among predicate and function symbols or between predicate and function symbols and variables: each predicate and function symbol in a sequence will take its value independently of the other predicate and function symbols: and each variable will take an object from the domain as a value, independently of the predicate and function symbols though not independent of the other variables.[12]

We now consider the rules for determining the semantic connections among more complex expressions. Let us begin with a couple of examples. Suppose we wish to determine

[12] We do not have to think of this rule as pertaining to the lexicon, since we might allow that an arbitrary sequence $E_1, E_2, \ldots, E_n$ of lexical items should be syntactically derivable from the subsequence of its variables and from each of the individual expressions in the sequence that are not variables. In general, we should be able to extract the ‘independent’ items when determining the semantic connection among various expressions.

the semantic connection, i.e. value range, on x + y. We should do this on the basis of the semantic connection on +, x, y. But a number *c* will belong to the value range of x + y just in case *c* is of the form $a + b$ for some triple ($+$, *a*, *b*) in the semantic connection on +, x, y. Note that it will not do to look at the respective values of +, x and y, since then we would be unable to distinguish between x + x and x + y. Suppose next that we wish to determine the semantic connection on x + y, x, y. We should do this on the basis of the semantic connection on +, x, y, x, y (with x + y in the sequence giving way to +, x, y). But the triple, (*c*, *d*, *e*) will belong to the semantic connection on x + y, x, y iff *c* is of the form $a + b$ for some quintuple ($+$, *a*, *b*, *d*, *e*) in the semantic connection on x + y, x, y, x, y. Note that it is here essential for the correctness of this result that we assume the Coordination Principle, since it is this that guarantees that $a = d$ and $b = e$ and hence that any (*c*, *d*, *e*) will belong to the semantic connection on x + y, x, y iff *c* is of the form $d + e$.

In general, we determine the semantic relationship between a complex expression and other expressions on the basis of the semantic relationship between the components of the complex expression and the other expressions. Let us use Δ and Γ for arbitrary sequences of expressions and δ and γ for arbitrary n-tuples of entities. Given that the expression E is complex, the general rules for determining the semantic connection [Δ, E, Γ] on Δ, E, Γ are as follows:

<u>for E a term of the form $ft_1t_2 \ldots t_n$</u>

$$[\Delta, ft_1t_2 \ldots t_n, \Gamma] = \{(\delta, f(e_1, e_2, \ldots, e_n), \gamma): (\delta, f, e_1, e_2, \ldots, e_n, \gamma) \in [\Delta, \mathrm{f}, t_1, t_2, \ldots, t_n, \Gamma]\};$$

<u>for E an atomic formula $Pt_1t_2 \ldots t_n$</u>

$$[\Delta, Pt_1t_2 \ldots t_n, \Gamma] = \{(\delta, \Pi(P, e_1, e_2, \ldots, e_n), \gamma): (\delta, P, e_1, e_2, \ldots, e_n, \gamma) \in [\Delta, \mathrm{P}, t_1, t_2, \ldots, t_n, \Gamma]\};$$

<u>E the negation ~F</u>

$$[\Delta, {\sim}F, \Gamma] = \{(\delta, -\pi, \gamma): (\delta, \pi, \gamma) \in [\Delta, F, \Gamma]\};$$

<u>E the disjunction $(F \vee G)$</u>

$$[\Delta, (F \vee G), \Gamma] = \{(\delta, \pi_1 \cup \pi_2, \gamma): (\delta, \pi_1, \pi_2, \gamma) \in [\Delta, F, G, \Gamma]\},$$

(We use some obvious notation: $\Pi(P, e_1, e_2, \ldots, e_n)$ is $\top$ or $\perp$ according as to whether or not $(e_1, e_2, \ldots, e_n) \in P$; - is the truth-functional operation of negation; and $\cup$ is the truth-functional operation of disjunction.)

If these rules are applied to a finite sequence of expressions, then they will eventually enable us to ascertain the semantic connection on the sequence on the basis of the semantic connection on a corresponding sequence of lexical items; the semantic connection on this latter sequence may then be ascertained according to the previously stated rules. As we have noted, there will be some ambiguity in how the rules might be applied; for when two complex expressions belong to a given sequence, we have a choice as to which is to be considered first. We could decide on a set order (e.g. left-to-right or right-to-left) but, even if we allow a choice, it is readily be shown that it makes no difference to the resulting semantic connection.

It should be clear that the relational semantics satisfies the various desiderata we have laid down and that is suffers from none of the defects of the previously considered approaches. The syntax is standard; the semantics is extensional; the connections are non-typographic; and compositionality and determination, in appropriate form, are secured. We can also directly explain the semantic relationship between two expressions E and F, since this is simply given by the semantic connection on E, F. We therefore see that, by rejecting semantical intrinsicalism and adopting a relational form of semantics, we are able to provide what appears to be a completely satisfactory semantics for the quantifier-free portion of predicate logic.

§4. The Problem of Bound Variables

We take up the question of extending the relational form of semantics to quantifiers (similar considerations will apply to other forms of variable-binding). The extension is not as straightforward as one might think; and, in the present section, we shall show that the most obvious proposals are subject to various difficulties while, in the following section, we show how these difficulties can be avoided if certain natural presupposition that we have so far made are abandoned. Again, this section is somewhat negative and the reader interested in my more positive proposals may more directly to §5.

The obvious way of extending our rules to quantified formulas is as follows:

for E the quantified formula $\forall$xF

$[\Delta, \forall xF, \Gamma] = \{(\delta, \pi, \gamma): \pi = \top$ if $(\delta, \top, a, \gamma) \in [\Delta, F, x, \Gamma]$ whenever $a \in D$, and $\pi = \bot$ otherwise$\}$.

Here $\forall$xF is taken to be derivable from the quantified formula F and the quantified variable x, as in the standard formu of syntax; and $\forall$xF is then taken to have the value $\top$ when F receives the value $\top$ for every assignment of an object to x compatibly with the assignment of values to the 'side' expressions.

However, there are two problems with this rule. The first is illustrated by the formula $\forall$x(x < y). Intuitively, $\bot$ is the only value in its value range. But (*a*, $\top$) will be an admissible assignment for x, x < y for any number *a*, since y may be assigned a number greater than *a*; and so, according to the rule, $\top$ should also be in the value range of $\forall$x(x < y). Clearly, what has gone wrong is that the rule is telling us that $\forall$x(x < y) is capable of assuming the value $\top$ if, for each assignment of a number to x, x < y is capable of assuming the value $\top$ for some assignment of a number to y. The value of y is being made to depend illegitimately on the value of x.[13]

I believe that this problem is largely a consequence of the fact that we are attempting to provide an extensional semantics. Suppose that we assign *propositions* to formulas rather than truth-values. Then from an admissible assignment (*a*, *p*) to x, x < y, say, we may determine an admissible assignment *q* for $\forall$x(x < y). For we may take *q* to be the proposition that results from *p* by removing *a* and then quantifying over the resulting propositional form. Thus the propositional value assigned to $\forall$xF(x) can be determined 'point-wise' from any one of the propositions that might be assigned to F(x). But the truth-value of $\forall$xF(x) can only be determined from the truth-value of *all* of its instances; and it is this circumstance that gives rises to the present difficulty.[14]

---

[13] A similar problem would arise if we were to formulate the rule for conjunction in any analogous manner:

$[\Delta, (F \wedge G), \Gamma] = \{(\delta, \pi_1 \cap \pi_2, \gamma):(\delta, \pi_1, \gamma) \in [\Delta, F, \Gamma] \text{ and } (\delta, \pi_2, \gamma) \in [\Delta, G, \Gamma]\}.$

For $\top$ is an admissible assignment for each of x > 0 and ~(x > 0) but not, of course, for (x > 0 $\wedge$ ~(x > 0)).

[14] Strictly speaking, the point-wise propositional rule only works if the object *a* in the 'x-position' of the proposition assigned to F(x) does not also occur elsewhere. We shall later see how a more careful formulation of the rule might be given.

Let us also note that we have here a solution to one of the difficulties in the instantial approach: for we may take $\forall$xF(x) to be syntactically derived from the generic instance F(x), not a particular instance F(c); and the propositional value of $\forall$xF(x) can then be determined on the

In any case, it may be resolved by insisting that the free variables of $\forall xF(x)$, or of other expressions, be 'pinned down'. Call a sequence of expressions $E_1, E_2, ..., E_n$ *full* if each variable free in one of its members is itself one of its members. Then the present difficulty in the rule for quantification disappears once we insist that the given sequence $\Delta, \forall xF, \Gamma$ be full.[15]

Since we need to be able to assign a semantic connection to every sequence of expressions, whether full or not, we now need a special rule to enable us to determine the semantic connection on non-full sequences. This takes the following general form:

Down

$$[E_1, ... , E_m] = \{(e_1, e_2, ..., e_m): (e_1, e_2, ..., e_m , e_{m+1}, ..., e_n) \in [E_1, ... E_m , ..., E_n]\}, n > m.$$

In other words, the connection on a given sequence of expressions may be obtained by restriction of the connection on a larger sequence.[16] Using the previous rules, we may first obtain the semantic connections on the full sequences; and then, by using the Down Rule, we may then obtain the semantic connections on the non-full sequences.

Of course, semantic connections on full sequences are analogous to the specification of semantic values under Tarski assignments. To say that (π, a, b) is an admissible assignment for $x > y$, x, y, for example, is to say that $x > y$ receives the truth-value π under the assignment of a to x and of b to y. Relational semantics is therefore coming to look very much like the Tarski semantics. But two points should be made about the resemblance. First, despite the similarity in formulation, our underlying semantic perspective is very different from Tarski's. For we do not take semantic values to be assigned to expressions relative to assignments, and nor do we take them to be the result of incorporating the assignments into the values themselves; the assignment of values to the variables and of semantic values to terms and formulas are treated strictly on a

---

basis of the *range* of propositional values of F(x). Thus there is no need to presuppose an understanding of any particular instance, as with the instantial approach.

[15]It might also be corrected by adopting an analogue of the original clause for conjunction in which each instance of the quantified formula is required to occur in the new sequence. But this involves considerable technical complications.

[16]For our purposes, we need only consider the case in which the additional expressions $E_{m+1}, ... E_n$ are the free variables of $E_1, ... E_m$. The more general rule gives rise to problems in the case of the empty domain.

par. Second, it was only rather special circumstances that led to the present convergence in formulation. We shall later come across other considerations that pull the two formulations apart.

I turn to the second problem, which involves having too *many* side variables, as opposed to too *few*. Consider the sequence $\forall$xPx, x. Then according to the quantificational rule, ($\perp$, *a*) will be an admissible assignment for this sequence only if ($\perp$, *b*, *a*) is an admissible assignment for the sequence Px, x, x for some object $b \in D$. But given coordination, this requires that $b = a$, i.e. that P be false of *a*. And clearly this is an incorrect result; if $\forall$xPx is the formula $\forall$x(x > 0), for example, then we want ($\perp$, 2) to be an admissible assignment for $\forall$x(x > 0), x, even though ($\perp$, *b*, 2) is never an admissible assignment for x > 0, x, x. The problem, of course, is that the quantified variable x may 'accidentally coincide' with some free occurrences of x in the rest of the sequence and hence may be inadvertently tied to their interpretation, once the quantifier is removed.

There are some obvious technical solutions to this problem, though none yields a philosophically satisfactory semantics. Perhaps the most obvious solution is to model the semantics more closely on the Tarski semantics. Call a sequence of expressions E, $E_1, \ldots, E_n$, $n \geq 0$, *exact* if E is a formula and $E_1, \ldots, E_n$ are *exactly* the free variables that occur in E. A exact sequence may be written in the form A, $x_1, x_2, \ldots, x_n$, where $x_1, x_2, \ldots, x_n$ are the free variables of A (repetitions are allowed). The admissibility of the assignment $(\pi, a_1, a_2, \ldots, a_n)$ for F, $x_1, x_2, \ldots, x_n$ under the relational semantics then corresponds to F receiving the truth-value π for the assignment of $a_1, a_2, \ldots, a_n$ to $x_1, x_2, \ldots, x_n$ under the Tarskian semantics. If we now restrict the quantificational rule above to exact sequences, we obtain the relational counterpart of the standard clause for universal quantified formulas under the Tarskian semantics:

<u>for E the quantified formula $\forall$xF and $x_1, x_2, \ldots, x_n$ the free variables of F</u>

$[\forall xF, x_1, x_2, \ldots, x_n] = \{(\pi, a_1, a_2, \ldots, a_n)$: $\pi = \top$ if $(\top, a, a_1, a_2, \ldots, a_n) \in [F, x, \Gamma]$ whenever $a \in D$, and $\pi = \perp$ otherwise$\}$.

Our problem now disappears, since the quantified variable x will not be among the variables $x_1, x_2, \ldots, x_n$.

This proposal is fine as far it goes, but it calls for further modifications in the semantics. For even if we start off with exact sequences, the induction will throw up sequences that are not exact; and since we should, in any case, assign a semantic connection to any sequence expression, we shall need add clauses to show how one is to evaluate the sequences that are not exact. We can use *Down* to evaluate the subsequences of exact sequences. But this still leaves us with the task of evaluating sequences that are not included in exact sequences. To this end, say that $E_{k1}, \ldots , E_{kp}$ is a *full-subsequence* of the sequence $E_1, \ldots , E_n$ if it is a subsequence that contains all variables of $E_1, \ldots , E_n$ that are free in one of $E_{k1}, \ldots , E_{kp}$ (in other words, if x is free in $E_{k1}, \ldots , E_{kp}$ and some $E_i = x$ then some $k_j = i$). We then adopt the following additional rule:

<u>Up</u> Suppose that $E_1, \ldots , E_n$ is composed of the full-subsequences $E_{k1}, \ldots , E_{kp}$ and $E_{l1}, \ldots , E_{lq}$ (they may overlap). Then $[E_1, \ldots , E_n] = \{(e_1, e_2, \ldots, e_n): (e_{k1}, e_{k2}, \ldots, e_{kp}) \in [E_{k1}, \ldots , E_{kp}]$ and $(e_{l1}, e_{k2}, \ldots, e_{lq}) \in [E_{l1}, \ldots , E_{lq}]\}$.

Whereas *Down* tells us how to evaluate downwards, from larger to smaller sequences, *Up* tells us how to evaluate upwards, from smaller to larger sequences.[17]

With the help of *Up* and *Down*, we are then in a position correctly to evaluate the semantic connection on *any* sequence of expressions.[18] But even though these rules produce the correct results, I do not believe that they produce them in the right way: the semantic evaluation of formulas does not accord with how they are understood. The problem is with *Up*. For consider how we are directed to evaluate the semantic connections on the respective sequences, $\forall$xPx, y and $\forall$yPy, y. The first will be evaluated on the basis of the connection on Px, x, y; $\forall$xPx will receive the value $\top$ when y takes a given value if and only if Px receives the value $\top$ whenever y takes the given value and x takes any value whatever. Thus in evaluating the truth-

[17] *Up* is an analogue of *Piecing* in Fine [85], p. 25. It is clear that the application of *Down* and *Up* will need to be restricted if we are to avoid evaluations that do not 'bottom out'.

[18] By *Down*, the semantic connection on a sequence of expressions can be determined on the basis of the semantic connection on a full sequence. So we may suppose that the given sequence of expressions is full. If the expressions in the sequence are all simple, then the semantic connection is given by the lexical semantics. So suppose that at least one of the expressions is not simple. The sequence is then of the form Δ, E, Γ for complex E. Let $x_1, x_2, \ldots, x_i, E, x_{i+1}, \ldots, x_n$ be an exact full-subsequence of Δ, E, Γ containing the given expression E. Then using *Up*, the semantic connection on Δ, E, Γ may be determined on the basis of the semantic connections on $x_1, x_2, \ldots, x_i, E, x_{i+1}, \ldots, x_n$ and Δ, Γ. Proceeding in this way, the evaluation will bottom out with connections on sequences of lexical items.

value of $\forall xPx$ for a given value of y, we keep the value for y fixed as we vary the value taken by x. But the evaluation of the second sequence proceeds quite differently: for we first evaluate the connection on $\forall yPy$ independently of y; we then evaluate the connection on y, independently of $\forall yPy$; and we then 'stick' the two evaluations together.

Although I have illustrated this difference in the case of sequences, it will also show up in the evaluation of formulas. Consider the formulas $(\forall xPx \wedge Qy)$ and $(\forall yPy \wedge Qy)$, for example. To evaluate the former, we must first combine it with the variable y (in the light of *Down*) and then evaluate $\forall xPx$, Qy, y while, to evaluate the latter, we must first combine it with y and then evaluate $\forall yPy$, Qy, y. Thus the very same kind of difference will emerge in the respective evaluations of $\forall xPx$, Qy, y and $\forall yPy$, Qy, y.

But these results are highly counter-intuitive. For we have a strong intuition that $(\forall xPx \wedge Qy)$ and $(\forall yPy \wedge Qy)$ are mere notational variants and that their semantic evaluation should proceed in the same way. Thus the semantics fails to respect what I previously called 'Process Correlation'. The underlying problem is in the syntax - not for the formulas, for which the syntax is standard, but for the sequences. For we feel that $(\forall xPx \wedge Qy)$ and $(\forall yPy \wedge Qy)$ should be subject to the same pattern of syntactic analysis (and it is on this basis that we judge that their semantic evaluation should be the same). But under the implicit syntactic analysis that provides the basis for the semantic evaluation, this is not so. For when, in the syntactic analysis of each formula, we reach the sequences $\forall xPx$, Qy, y and $\forall yPy$, Qy, y, the two syntactic analyses diverge, with the first reducing to Px, x, Qy, y and with the second being decomposed into $\forall yPy$ and Qy, y. We therefore see that the present semantics inherits some of the typographic bias of the original Tarskian semantics, though in a much more subtle form.

How is this last residue of typographic bias to be removed? Let us briefly consider two variants of the present proposal, though neither will ultimately prove to be successful. According to the first, we should attempt to model the quantifier rule more closely upon the corresponding rule in the Tarski semantics:

<u>where E is the quantified formula $\forall xF$ and $x_1, x_2, ..., x_n$ include the free variables of F</u>

$[\forall xF, x_1, x_2, ..., x_n] = \{(\pi, a_1, a_2, ..., a_n): \pi = \top$ if $(\top, a, b_1, b_2, ..., b_n) \in [F, x, x_1, x_2, ..., x_n]$ whenever $a \in D$ and the assignment of $b_1, ... , b_n$ to $x_1, x_2, ... , x_n$ is an x-variant of the assignment

of $a_1, ..., a_n$ to $x_1, x_2, ... , x_n$, and $\pi = \perp$ otherwise}.
However, as it stands, this rule will not do, since the condition of being an x-variant introduces typographic bias into the truth-conditions. We must somehow express the same idea without reference to the variable. This might be done by appealing instead to all those positions j, $1 \leq j \leq n$, for which $c_j = c$ whenever $(\pi, c, c_1, ..., c_n)$ is an admissible assignment for F, x, $x_1, x_2, ..., x_n$. It is then the b_j in $b_1, ... , b_n$, for these positions j, which may differ from the a_j in $a_1, ..., a_n$.

The resulting clause is exceedingly artificial. But it also gives rise to a subtle form of the previous problem. For although we have a single rule to cover the cases in which the side variables $x_1, x_2, ..., x_n$ contain the variable x and the cases in which they do not, these two cases are not treated in a uniform manner. Consider again how we are to evaluate the sequences ∀xPx, y and ∀yPy, y. In the first case, we proceed by evaluating Px, x, y and, in the second case, by evaluating Py, y, y. But semantic connections on these sequences will be different. So again, we have an unwanted divergence in evaluation, though one that now manifests itself, not as a difference in the syntactic analysis of corresponding quantificational sequences, but as a difference in the semantic connections that are to be considered once the corresponding syntactic analyses are made.

A rather different tack we might take is to re-letter formulas so that the offending instances of the quantificational rule never arise. Call a sequence of expressions $E_1, ... , E_n$ *differentiated* if (i) no variable occurs both free and bound in the sequence and (ii) no quantifier expression ∀x occurs twice in the sequence (for any given variable x). Given any formula A (or sequence of expressions $E_1, ... , E_n$), we may rewrite its bound variables so as to obtain a logically equivalent differentiated formula (or sequence). The undifferentiated formula $(\forall xPx \wedge \forall xQx) \wedge Px$, for example, may be rewritten as $(\forall yPy \wedge \forall zQz) \wedge Px$. This suggests that we might restrict the application of the quantificational rule to differentiated sequences, without any loss of expressive power, and thereby avoid having to deal with accidental coincidents.

This proposal might be justified if the use of variables in undifferentiated formulas were genuinely ambiguous. For we might then regard the differentiation of a formula as a form of disambiguation; and the proposal would merely amount to the requirement that a formula be disambiguated prior to its semantical evaluation. But it is hard to accept that the use of variables in undifferentiated formulas is - or, at least, must be - ambiguous. Consider the use of the variable x in the formula $\forall xPx \wedge \forall xQx$, for example. Surely it must be possible to use the variable x with the same 'meaning' throughout this formula; and if we do, then what reason could we have for thinking that its interpretation must be different from what we normally take it to be? The problem of accidental coincidents is therefore still with us.[19]

§5. Relational Semantics for Bound Variables

I shall now propose what I regard as a more satisfactory solution to the problem of accidental coincidents. We shall see that the solution requires a fundamental revision in the framework of syntactic and semantic analysis that worked so well in the case of free variables. For we can no longer tie coordination to the identity of variables; we must allow cases in which there is no coordination even though the variables are the same.

Our problem was to find a uniform way of evaluating quantified formulas, one that pays no regard to 'accidental' coincidences among occurrences of the same variable. In particular, the sequences $\forall xPx$, y and $\forall yPy$, y should, but for their notational difference, be evaluated in the very same way. If we adopt the standard quantificational rule, then the former is evaluated on the basis of Px, x, y and the latter on the basis of Py, y, y. But given that the last two variables in Px, y, y are coordinated while the last two variables in Px, x, y are not (assuming, of course, that there are at least two objects in the domain), it follows that the semantic connections on these sequences will not be the same and hence that the evaluation cannot be uniform. We considered various more or less ingenious ways of trying to restore uniformity but all of them appeared to flounder on the same or similar difficulties.

I would now like to propose a solution by 'brute force'. In the auxiliary evaluation of Px, x, y, the last two variables should be independent. Uniformity therefore requires that, in the auxiliary evaluation of Py, y, y, the last two variables (or variable occurrences) should also be independent. I therefore propose that we take these variables to be independent, even though

[19] It is a common technique to specify a semantics for a language by specifying a semantics for a language into which it is translated. But, as this case shows, it may not always be clear how to treat the product of a translation and a semantics as another semantics.

they are the same!

What prevented us from adopting this proposal before was the implicit assumption that different (free) occurrences of the same variable should be coordinated. This must now be given up. But how can this be on the cards? For was not coordination a centerpiece of the relational approach? And was it not through coordination - and coordination alone - that relationism was seen to be necessary and taken to be manifest? Our intention, however, is not to give up coordination altogether, but merely its indiscriminate application to every case in which we confront different occurrences of the same variable. Benign applications - however they are to be distinguished - should be retained, and the bothersome ones thrown out.

Simple as this idea may be, its development calls for some fundamental revisions in the syntax and the semantics of our current approach. Let us focus on the syntax since, once the syntax is in place, the semantics should follow. We then immediately face a problem in characterizing the objects of evaluation. If Py, y, y is an initial sequence, the last two variables should be coordinated but, if it is plays an intermediate role in the evaluation of $\forall y Py$, y, the last two variables should *not* be coordinated. But there is nothing in the sequence itself to indicate whether it should be regarded as initial or intermediate. We must therefore explicitly indicate whether or not the variables are to be coordinated and this additional information should always be supplied when we wish to consider how to evaluate a given sequence of expressions.

So suppose that we are given a sequence of expressions $E_1, \ldots, E_n$. Let us call an equivalence relation $\mathscr{C}$ on the (free and bound) occurrences of variables in $E_1, \ldots, E_n$ a *coordination-scheme* for the sequence.[20] Thus a coordination-scheme for Py, y, y might relate all three occurrences of the variable y to one another or it might relate the first two only or the second two only or relate no two distinct occurrences at all. Given a sequence of expressions $\Delta = E_1, \ldots, E_n$ and a coordination-scheme $\mathscr{C}$ for Δ, we might call the ordered pair $\Delta; \mathscr{C}$ a (*coordinated*) *system of* expressions. Thus a coordinated system of expressions contains an in-

[20]I shall talk loosely and informally of occurrences, not always taking care to distinguish between an expression and its occurrences. There are interesting issues, which I shall not pursue, as to how exactly occurrences and the various key operations and relations on them should be defined.

built scheme of coordination; and it is these systems, rather than the unadorned sequences, that should now be taken to be the proper object of semantical study.[21]

It should be observed that coordination is, in a certain sense, a purely formal relation. Given the identical sequence of variables x, x there is no *fact* as to whether or not its two occurrences of x are coordinated. Rather, we are free to stipulate whether or not the relation is to hold between them, thereby obtaining either a coordinated or an uncoordinated pair of variables. Thus the relation is one by which syntactic objects are *defined* rather than *described*. The relation does have its empirical side though ; for whether two successive *tokens* of the variable x are tokens of a coordinated or uncoordinated sequence will depend upon whether the tokener *intends* them to be tokens of a coordinated or uncoordinated sequence (which, in its turn, will be determinative of whether he intends them to be semantically coordinated).[22]

If we treat coordination as a syntactic relation, that serves to underlie a corresponding semantic relation, then we see that whether or not this relation holds between two occurrences of variables is not something that will in general supervene on the intrinsic syntactic features of the variables themselves. Given the variables x, x there is nothing in the variables themselves to say whether or not they are to be coordinated. We are therefore led to embrace a form of syntactic relationism that is the direct counterpart of the form of semantic relationism that we have already felt obliged to endorse.

Although we are free to stipulate whether or not two variables are coordinated, some schemes of coordination are more natural than others. It seems reasonable, in particular, to lay down the following requirements on coordination (taking for granted our previous requirement that $\mathscr{C}$ should be an equivalence relation on all occurrences of variables in the sequence Δ):

---

[21]Alternatively (not that this is the option I prefer), one might suppose that sequences of expressions are still the proper object of semantical study but that semantic values are now assigned to them *relative* to a scheme of coordination.

[22] Let me not attempt to enter into a deeper metaphysical explanation of what formal relations are. I suspect that they are best viewed as a kind of ‘postulational primitive’ in the sense of Fine [03].

(i) The F-F Rule Two free occurrences of variables are to be coordinated iff they are occurrences of the same variable;

(ii) The B-B Rule Two bound occurrences of variables are to be coordinated iff they are bound by the same quantifier;[23]

(iii) The B-F Rule No bound occurrence of a variable should be coordinated with a free occurrence of a variable.

Let us call these requirements the *default* rules; and let us call any coordination scheme $\mathscr{C}$ for Δ that conforms to them a *default* scheme and any system with a default scheme a *default* system. It should then be clear that the default scheme for any sequence of expressions is unique and corresponds to the way in which its expressions are most naturally understood.

Although the default is natural, it need not be adopted. Suppose, for example, that I have a defective typewriter: sometimes only the x-key work, and sometimes only the y-key. I may then stipulate that x and y are to be treated as mere typographic variants, with the consequence that x and y will be coordinated in Rxy, for example, or in ∀xPy. This is certainly confusing to the eye but it is does not constitute an illegitimate use of variables (and nor. I am inclined to think, does it have the effect of making 'x' and 'y' the same symbol). Or again, we may use the same symbol for 'pawn' to signify different pawns in representing the position of pieces in a game of chess. We find it natural to interpret variables by means of the default; but the default is not intrinsic to our very understanding of what it is to be a variable.

One place where the default must be given up, ironically enough, is in developing the semantics for predicate logic! For, as we have seen, even if we insist that the formulas of the symbolism be subject to the default, a properly defined induction for the semantics will require us to consider systems that are not subject to the default. Indeed, it was essentially our earlier insistence that semantical evaluation be confined to default systems that led to the difficulties over achieving uniformity. Thus dropping the default has become, for us, a theoretical necessity and not merely an idle curiosity.

[23]The occurrence of a variable x in a formula A is said to be *bound by* the occurrence of the quantifier ∀y in A if x = y and the occurrence of x is contained in the occurrence of the subformula ∀yB of A that begins with the occurrence of ∀y.

Which relaxations of the default are required to carry through the proposed revision to the semantics? Perhaps the weakest relaxation we might make is to replace the F-F rule with:

(i)′ The Weak F-F Rule Two free occurrences of variables are to be coordinated only if they are occurrences of the same variable.

Thus free occurrences of the same variable may now be uncoordinated; and so we are free to ‘decompose’ $\forall x F$ into B and x within a sequence without ‘accidentally’ coordinating the given occurrence of x with any other free occurrences of x in the sequence.

It seems to me, however, that further weakening of the constraints is possible consistent with our general semantic intentions. Indeed, we can drop the first two constraints altogether and simply insist upon a strengthening of the B-F rule:

(iii)′ The Strong B-F Rule No occurrence of a variable x that belongs to a quantifier $\forall x$ should be coordinated with a variable that does not fall within the scope of the quantifier.

This new constraint is quite weak: it never requires us to coordinate two variables; and it allows us to coordinate any two variables, even distinct variables, as long as their coordination does not result in a case of ‘quantifying out’.[24]

Even this constraint might be dropped, though the resulting semantics will be anomalous in various ways. Consider the formula $(\forall xPx \wedge \exists xQx)$, for example, and let us suppose that all four occurrences of x are taken to be coordinated. We evaluate $(\forall xPx \wedge \exists xQx)$ in terms of the sequence $\forall xPx$, $\exists xQx$. Now in evaluating this sequence, we can either work first on $\forall xPx$ or work first on $\exists xQx$. But under the former, the formula will be equivalent to $\forall x(Px \wedge Qx)$ while, under the latter, it will be equivalent to $\perp$ (given that there are at least two objects in the domain). Thus the order of evaluation will make a difference to the truth-conditions. Of course, one could insist that the first conjunct always be evaluated first, but then $(\forall xPx \wedge \exists xQx)$ would not be

[24] Although the constraint is very liberal, it does not result in any increase in expressive power; for, given any coordinated formula subject to the constraint, we may rewrite the occurrences of variables in A in such a way that the corresponding coordination scheme is the default on the resulting formula A′. Conversely, given a coordinated formula A not subject to the constraint, there will be no way to rewrite its variables so that the resulting formula A′ is subject to the default coordination.

equivalent to ($\exists$xQx $\wedge$ $\forall$xPx).[25]

With the syntax in place, let us now consider how the semantics should go. The principal modification is to the Coordination Principle, which must now be 'sensitized' to the various relationships of coordination that may hold among the variables:

Coordination (on Systems) $[x_1, x_2, ..., x_m; \mathscr{C}] = \{(a_1, a_2, ..., a_m): a_1, a_2, ..., a_m \in D$ and $a_i = a_j$ if x_i is coordinated (through $\mathscr{C}$) with x_j, $1 \leq i, j \leq m\}$.

Thus coordination is now constrained by the coordinating relation $\mathscr{C}$ rather than by identity.

Given this syntax, the semantics proceeds much as before, but with systems of expressions in place of sequences of expressions. For example, the semantical rule for negative formulas now take the following form:

$[\Delta, \sim F, \Gamma; \mathscr{C}] = \{(\delta, -\pi, \gamma): (\delta, \pi, \gamma) \in [\Delta, F, \Gamma; \mathscr{C}']\}$;

where $\mathscr{C}'$ is the scheme of coordination on Δ, F, Γ that corresponds to the scheme of coordination $\mathscr{C}$ on Δ, $\sim$F, Γ.[26]

On the other hand, the rule for quantified formulas becomes:

$[\Delta, \forall xF, \Gamma; \mathscr{C}] = \{(\delta, \pi, \gamma): \pi = \top$ if $(\delta, \top, a, \gamma) \in [\Delta, F, x, \Gamma; \mathscr{C}']$ whenever $a \in D$, and $\pi = \perp$ otherwise$\}$,

where Δ, $\forall$xF, Γ is (in the appropriate sense) full and $\mathscr{C}'$ is again a scheme of coordination for Δ, F, x, Γ that corresponds to the scheme of coordination $\mathscr{C}$ or Δ, $\forall$xF, Γ. Note that the 'isolated' x in Δ, F, x, Γ; $\mathscr{C}'$ will stand in the 'same' coordinative relationships as the bound x in Δ, $\forall$xF, Γ; $\mathscr{C}$, and so unwanted links to other occurrences of x in Δ, Γ will be avoided. Since the quantificational rule is restricted to full systems, we shall need the appropriate version of *Down*,

---

[25] Note that, when y is coordinated with x, ($\exists$xPx $\wedge$ Qy) is equivalent to $\exists$x(Px $\wedge$ Qx), which is the result one would expect under the 'dynamic' semantics for predicate logic. We see from this example how something like the present framework might be used to provide a common rubric for the classical and dynamic interpretations of predicate logic.

[26] There is a tricky question as to how exactly these relationships on coordination schemes are to be defined. Probably the simplest technical solution is to suppose that no variable is allowed to occur twice in any sequence of expressions. There is then no need to distinguish between a variable and its occurrences and so $\mathscr{C}'$, for example, can simply be identified with $\mathscr{C}$.

but the problematic *Up* rule will not be required.

The statement of the semantics is now complete; and I hope it is clear that it as satisfactory as the semantics we developed for free variables (modulo the need to appeal to *Down*). The basic device that has enabled the semantics to work is the accommodation of semantical relations. Thus just as the work of Tarski's indicated the need to extend the notion of *truth* to *satisfaction* in order to achieve a technically adequate semantics for predicate logic, so the present work indicates the need to extend the notion of semantic *value* to semantic *connection* in order to achieve a philosophically adequate semantics. With words as with people: only connect.

§6. The Character of Coordination

I wish to draw out various general implications and morals for the syntax and semantics of coordination from our previous discussion.

We should observe, in the first place, that the current semantics avoids the earlier difficulties over accidental coincidents. When we evaluate $\forall xPx$, x; $\mathscr{C}$ under the default scheme using the intermediary Px, x, x; $\mathscr{C}'$, the first two occurrences of x will not be coordinated with the others and so the rule will not yield unwanted results. The evaluation of $\forall xPx$, x under the default scheme will also be 'uniform' with the evaluation of $\forall xPx$, y under a default scheme, since the relevant patterns of coordination in Px, x, x and Px, x, y are exactly the same.

A more general result holds, which I state for formulas, though it also holds for sequences. Suppose that the formula A′ is obtained by rewriting the free and bound occurrences of variables in A and that the coordination-scheme $\mathscr{C}'$ for A′ corresponds, under this rewriting, to the coordination-scheme $\mathscr{C}$ for A. Then A′ and A, under their respective schemes of coordination, will be evaluated in exactly the same way; their 'process' values will be the same. Thus the kind of uniformity in semantical evaluation that we found so hard to achieve before is here immediately forthcoming.

A particular case of this result is of special interest. For in rewriting A as A′, we may let all of the variables be the same. For example, $\forall x\forall yRxy$ (with default coordination) could be rewritten as $\forall x\forall xRxx$ with coordination between the first and third and the second and fourth occurrences of the variable x. In this way, the need to use different variables could be avoided. We would therefore have a kind of vindication of the view (suggested in Russell [03], §§88, 93) that there is only a single variable for any given range of values. What would normally show up as differences in the variable would now show up as differences in the pattern of coordination.

The single variable case is interesting in another way. For coordinated formulas written with a single variable correspond to formulas written in what one might call ‘telegraphic notation’. This is the notation in which variables are replaced by blanks and the blanks are appropriately connected by lines. The blanks then correspond to the different occurrences of a single variable and the lines to the relationships of coordination. A number of philosophers have thought that the telegraphic notation provides an especially perspicuous way of representing the role of variables. It is therefore surprising, and also unfortunate, that they did not attempt to make the syntax and semantics for this notation precise, since they would then have arrived at something like the present, more philosophically informed conception of what the syntax and semantics of predicate logic should be.

The coordinative approach also make clear that the introduction of quantification or binding requires the use of essentially new ideas: the semantics for bound variables is not simply implicit in the semantics for free variables. We could imagine a linguistic novice who had a perfect understanding of the role of variables and of the meaning of quantifier phrases and yet was unsure how to interpret quantified formulas. We might illustrate the nature of his predicament with the formula $(\exists xPx \wedge Qx)$. For it would not be evident to him whether he should take the first occurrence of x to bind the third (which is perhaps why beginners to logic have such difficulty in appreciating that the third occurrence of x is meant to be ‘dangling’). Even if we take it to be a general part of his understanding of free variables that only occurrences of *identical* variables can be coordinated, he would still have no basis for deciding whether the first occurrence of x was to coordinated to the third occurrence. And even if he could somehow gather that one was not allowed to quantify out, it would still be unclear to him whether the formula $\exists x\forall xPx$ should be interpreted as being equivalent to $\exists xPx$, for example, rather than to $\forall xPx$. Thus his understanding of free variables and of quantifier phrases must be supplemented by a more detailed understanding of the conventions governing coordination, or ‘binding’.

I consider it a great virtue of the present approach that it provides the kind of flexibility that enables one to see the semantics of predicate logic as a special case of a more general semantical scheme. The approach isolates 'parameters' in terms of which our use of variables should be given; and any particular understanding of variables can then be taken to consist in our fixing on certain 'values' of these parameters as opposed to others. Other approaches merely isolate one particular way in which the values of the parameters might be specified and so provide no real explanation of the kind of interpretative decision that should be seen to lie behind the adoption of one use of the variable as opposed to another.

A question we have not so far considered is how we should situate the current coordinative syntax and semantics for predicate logic within a general syntactico-semantic framework. This is a large and difficult question and all I can do is throw out a few hints and suggestions. Consider first the general syntax appropriate for 'sequences'. As we have already noted, one sequence may be derived from another on the basis of an underlying syntactic analysis of one of its components. Thus, given that ~F is derivable from F we make take Δ, ~F, Γ to be derivable from Δ, F, Γ. Call such derivations *induced*. But we may also allow a sequence to be derived from others in a manner that is *not* induced by an underlying syntactic analysis of one of its components. We might call such a derivation *sui generis* rather than 'induced', since it is peculiar to the sequences themselves. Thus under the approaches we considered above, the application of *Down* presupposed a sui generis derivation of smaller sequences from larger sequences while the application of *Up* presupposed, conversely, the derivation of a larger sequence from smaller sequences. I do not know what a full set of reasonable restrictions on the sui generis derivations might be, though their addition to the induced rules should certainly allow that every sequence of expressions containing a non-lexical item still be ultimately derivable from a sequence of lexical items.

I consider it a kind of ideal that there be no sui generis derivation since it can only be because of some restriction on the application of the induced derivations that their use is required. We saw that the semantics of free variables had no need of sui generis derivations, but that the semantics of bound variables did - either through the use of *Down* or (on certain versions of the semantics) of *Up*. But the only reason we had any need of these further rules is that the application of the quantificational rule had to be restricted in some way. As I have indicated, I think this 'blot' on our semantics is a consequence of the attempt to be extensional and would disappear under a suitable intensional form of the semantics.

Consider now the general syntax appropriate for systems. In the general case, a coordinative syntax might allow there to be several syntactic relations of coordination, with each restricted as to the syntactic category of the expressions between which it may hold. We may define a coordinated system in the same way as before, but with several coordinative relations in place of one.[27] The induced rules of derivation on sequences can be extended to systems in a straightforward and uniform way. Consider, for example, the rule that allows us to derive the sequence Δ, ~F, Γ 'negatively' from Δ, F, Γ. The corresponding rule on systems will allow us to derive Δ, ~F, Γ;$\mathscr{C}$ 'negatively' from Δ, F, Γ;$\mathscr{C}'$, where $\mathscr{C}'$ is the coordination scheme on Δ, F, Γ corresponding to the coordination scheme $\mathscr{C}$ on Δ, ~F, Γ. [28] I do not know if there is a reasonable uniform procedure for extending sui generis rules from sequences to systems, but we might note that the rules of derivation underlying *Up* and *Down* can both be extended to systems in the obvious way.

It is somewhat harder to say how the general semantics of sequences and systems should go, since it is not clear how far and in which direction we should generalize the current semantics. If e is the semantic value of E and f of F, then we may take the semantic connection on the sequence E, F, in the absence of coordination, to be a suitable 'product' $e \times f$ of e and f (for

[27]There is something to be said for thinking of a system more generally as a relational structure whose domain consists of the expressions and whose relations are the various coordinative relations along, perhaps, with an ordering.

[28] We need to say in general how the one scheme of coordination is determined from the other. So suppose that E is derived (in some manner) from $E_1, \ldots, E_n$. We then wish to say which occurrences of symbols in E are aligned with which occurrence of symbols in $E_1, \ldots, E_n$. Now $E_1, \ldots, E_n$ will be a substitution-instance of some expressions $F_1, \ldots, F_n$ which contain each symbol at most once, i.e. there will be a substitution-function s, taking symbols into symbols, which is such that each $E_i = s(F_i)$ for i = 1, ... , n. We suppose that there is an expression F which is derived from $F_1, \ldots, F_n$ in the same manner in which E is derived from $E_1, \ldots, E_n$ and that E = s(F). An occurrence of a symbol in E will then be aligned with the occurrence of a symbol in $E_1, \ldots, E_n$ if each results from applying the substitution-function s to the very same symbol (in F or $F_1, \ldots, F_n$ respectively).

us, × is the Cartesian product so that, if the semantic value of each of the variables x and y is the domain D, then the semantic connection on x, y will be $D \times D$). The semantic value of a syntactic relation R of coordination might be taken to be an appropriate *restriction* r. Thus if R relates the expressions E and F with semantic values e and f, then the semantic connection on E, F, as coordinated by R, will be the restriction $r(e \times f)$ of $e \times f$ (for us, the semantic value of the relation of coordination is the function taking a relation e into $e \cap \iota_D$, where ι_D is the identity relation on the domain D). Of course, 'separated' expressions, such as E_1 and E_3, might be coordinated within a sequence E_1, E_2, E_3 and, in this case, one might appeal to general algebraic operations on products to define the semantic connection. Thus if f $(= r(e_1 \times e_3) \times e_2)$ is the semantic connection on the correspondingly coordinated sequence E_1, E_3, E_2, then the semantic connection on E_1, E_2, E_3 might be obtained by 'permuting' the second and third arguments in f.[29]

I find it extraordinary, given the central role of the symbolism of predicate logic, that there currently exists no more general standpoint from which we might come to a better understanding of its syntax and semantics.

Let me conclude by addressing a general concern that the reader may have had. Under our approach, the syntax and semantics of predicate logic is much more complicated than under Tarski's. Is this not a mark against it? The answer is: only if it is attempting the same thing, but it is not. Tarski's project is best regarded as an attempt to provide a *technically* adequate semantics for the language of predicate of logic, while ours is best seen as an attempt to provide a *philosophically* adequate semantics, one that does not merely deliver a correct definition of truth but does so in such a way as to meet certain other desiderata, such as Determination or Neutrality or the preservation of meaning under notational variance. Of course, the complexity of our approach would be a mark against it if there an easier way to achieve the same goal. But I doubt that there is.

In this connection, it is worth bearing in mind that there is, in general, a trade-off between Neutrality and Simplicity. As I have noted, if compositionality were all that mattered, then it

[29] Ironically, this is a place within the semantics for predicate logic where we need something like relational algebra! Let me note that, when there are several coordinating relations, it may matter in which order they are applied.

might be achieved by taking each constituent expression to be its own semantic value; the semantics of any language would then be as simple as its syntax. Tarski insists on an adequate semantics for closed sentences, the semantic value of each closed sentence should be its truth-value; and this complicates how the semantics might be achieved. We insist on an adequate semantics for sentences and terms, be they open or closed, that is faithful to the way in which they are actually understood; and one would expect this to complicate the semantics even further.

It is by insisting that the semantics be at an appropriate distance from the language that we achieve a deeper understanding of what the semantics actually is. One would be surprised if an account of the semantics of variables were to pay no attention to their range of values (as under the 'identity' semantics mentioned above). Clearly, the possession of a range of values is a basic feature of variables and should be integral to any reasonable account of their representational role. Similarly, it seems to me, for the phenomenon of coordination, the fact that some variables take their values independently of or in coordination with others; this too should be integral to any reasonable account. Under the Tarski semantics, however, coordination, or 'binding', is merely an afterthought and plays no part in the semantics itself. Under our own approach, by contrast, the desire to achieve neutrality actually forces us to incorporate the phenomenon of coordination into the semantics.

I think that we have been guilty of a certain philosophical complacency over the semantics of predicate logic. The symbolism is so familiar and the Tarski semantics so perspicuous that it is hard to believe that we do not yet possess a full theoretical understanding of the way variables are used. But I hope to have shown that this understanding is something of an illusion: it does nothing to assuage the kind of fundamental worries concerning the use of variables that Russell voiced almost a century ago; and it is only through making radical revisions in the Tarski semantics that we are able to deal with these difficulties and achieve the kind of theoretical understanding that we can reasonably expect to have.