CHAPTER 3

FIRST-ORDER PREDICATE LOGIC

I. THE LANGUAGE L_1

We turn now to the construction of a formal language L_1, which adds individual variables and quantifiers over these variables to the syntactic apparatus of the language L_0. In fact, this language L_1 *contains* L_0, in the sense that all the sentences of L_0 will also be sentences of L_1 (but not conversely), and these sentences of L_0 will have to be interpreted just as before. As these new individual variables (for which we will use the symbols v_1, v_2, v_3, . . .) behave syntactically just like individual constants, we will introduce the new syntactic category of *individual terms* (or simply *terms*) to include both variables and constants.

1. *Syntax of* L_1

A. The *basic expressions* of L_1 are of four categories:
   1. the *names* d, n, j, and m
   2. a denumerably infinite supply of *individual variables* v_1, v_2, v_3, . . .
      (To avoid having too many subscripts in our formulas we will sometimes use x to stand for v_1, y to stand for v_2, and z to stand for v_3.) Together, the names and individual variables of L_1 comprise the *terms* of L_1.
   3. the *one-place predicates* M and B.
   4. the *two-place predicates* K and L.

B. The formation rules of L_1 consist of the following:
   1. If δ is a one-place predicate and α is a term, then $\delta(\alpha)$ is a formula.
   2. If γ is a two-place predicate and α and β are terms, then $\gamma(\alpha, \beta)$ is a formula.
      If ϕ and ψ are formulas, then so are:
   3. $\neg\phi$
   4. $[\phi \wedge \psi]$
   5. $[\phi \vee \psi]$
   6. $[\phi \rightarrow \psi]$
   7. $[\phi \leftrightarrow \psi]$

8. If ϕ is a formula, and u is a variable, then $\forall u\phi$ is a formula.
9. If ϕ is a formula, and u is a variable, then $\exists u\phi$ is a formula.

(Note that in clauses (8) and (9) u is used as a meta-language variable ranging over the variables of L_1; this is the only case of a non-Greek letter being used as a meta-language variable.)

In these rules we have used the term 'formula' where the term 'sentence' was used earlier. This follows the traditional practice of reserving the term *sentence* for a formula containing no free occurrences of variables, whereas formulas may or may not contain free occurrences of variables. The distinction between *free* and *bound* occurrences of variables is no doubt already familiar to the reader: an occurrence of a variable u in a formula ϕ can be defined as *bound in* ϕ if it occurs in ϕ within a sub-formula of the form $\forall u\psi$ or $\exists u\psi$; otherwise, that occurrence is *free* in ϕ. The syntactic rules do nothing to avoid so-called *vacuous quantification*, quantification over a formula with respect to a variable that does not occur in it; e.g. we can form $\forall xK(j, m)$ from $K(j, m)$, or $\exists yB(x)$ from $B(x)$. It would complicate the syntax greatly to prohibit such formulas, and moreover, the semantic rules will turn out to treat such vacuously quantified formulas as if the vacuous quantifier simply weren't there: $\forall xK(j, m)$ will be interpreted exactly like $K(j, m)$, and $\exists yB(x)$ will be interpreted like $B(x)$. Hence the vacuous quantifiers are harmless if we are willing to ignore them.

The novelty in the interpretation of L_1 over L_0 lies primarily in the notion *satisfaction of a formula by an assignment of objects to variables* (which is Tarski's term) or, to use the completely equivalent notion found in PTQ, the *truth of a formula with respect to an assignment of values to variables*.

The need for such a notion within a compositional semantics can perhaps be best grasped intuitively in the following way. In introductory treatments of elementary logic it is usually said that formulas with free variables cannot be either true or false as they stand because the variables themselves do not denote any particular individuals, hence the formulas make no real assertion until the variables have been quantified (or perhaps replaced with names in the course of a deduction). Thus we have Russell's term *propositional function* for such formulas; they are regarded not as propositions but as functions which give propositions when supplied with individuals as arguments. Quantified expressions are then treated in the following way: a sentence $\forall xB(x)$ is true just in case $B(x)$ is always true when x is regarded as denoting any member of the domain of discourse whatsoever; similarly, $\exists xB(x)$ is true if there is at least one individual in the domain such that

$B(x)$ is true when x is regarded as denoting that individual. (Alternatively, $\forall x B(x)$ may be said to be true when the result of substituting for x in $B(x)$ a name of any individual in the domain of discourse always gives a true sentence. This approach leads to the *substitutional theory of quantification*, but we will not be concerned with that theory here.) The treatment of multiple quantifiers and quantifiers binding arbitrarily complex sentences is then left pretty much to intuition, except for pointing out that it is necessary to distinguish scope relations carefully. For example, it might be said that to find out whether $\forall x \exists y L(x, y)$ is true we must "try out" all values for x, and for each one of these values of x we must try to find some value for y that makes $L(x, y)$ true. On the other hand, for $\exists y \forall x L(x, y)$ we must find a single value for y according to which $L(x, y)$ will remain true as the value of x is allowed to range over every member of the domain in turn. Such procedures are usually not rigorously specified, though of course explicit procedures may be given for carrying out *deductions* involving multiple quantifiers and arbitrarily complex formulas.

A little thought should convince the reader that something new will be needed if we are to give explicit truth conditions for quantified sentences by the compositional principle followed for L_0. Since the rules introducing quantifiers (B.8 and B.9) make a sentence $\forall u \phi$ from any sentence ϕ and any variable u, we must be able to give a completely general semantic rule for the truth conditions for $\forall u \phi$ in terms of the truth or falsity of ϕ, no matter what ϕ is. Now consider what will happen when we try to determine the truth conditions for $\forall x \exists y L(x, y)$ by this method. This formula is syntactically formed by first constructing the atomic sentence $L(x, y)$, then adding a quantifier by B.9 to give $\exists y L(x, y)$, then finally adding a second quantifier by B.8 to give $\forall x \exists y L(x, y)$. The semantic rules will retrace these stages. Thus the truth definitions corresponding to syntactic rule B.9 will have to give the truth condition for $\exists y L(x, y)$ in terms of $L(x, y)$ (and later, the rule corresponding to B.8 will have to give the conditions for $\forall x \exists y L(x, y)$ in terms of $\exists y L(x, y)$). But at the stage where y is quantified, the formula $L(x, y)$ has, in addition to the variable y being quantified at this step, the variable x which is still free as far as B.9 is concerned. Note also that in some syntactic derivations involving $\exists y L(x, y)$ the variable x will later be bound by a universal quantifier, and in other derivations by an existential quantifier. Of course this is the minimally simple case of multiple quantifiers; other cases will require that the semantic rule corresponding to B.8 give truth conditions for $\forall u \phi$ where ϕ has besides u any number of other free variables which will be bound at later stages. Obviously, free variables have to be dealt with *as such* by the semantic rules if a compositional semantics is to be given. As Tarski put it, "compound sentences are constructed from simpler sentential functions [i.e., formulas that may have free variables], but not always from simpler sentences" (Tarski 1944, p. 353).

It seems that what we need for these cases is some systematic means of "pretending" that each free variable denotes some individual or other, and then later systematically revising our assumption about which individual is denoted by these variables as we reach the appropriate quantifier at the "outer" or "higher" stages of the syntactic formation of the formula. Tarski's notion of *satisfaction* (or *truth relative to an assignment to variables*) is designed to do just that. Accordingly, we will add to the semantic machinery already described for L_0 the notion of a function assigning to each variable of L_1 some value from the domain A. Such a function is called an *assignment of values to variables*, or simply a *value assignment*. We will use the symbol g to denote such a function. (Since we have given L_1 an infinite supply of variables, the domain of g will be infinite, but functions with infinite domains present no particular problems.) The function g need not be one-to-one of course; in fact, it is important that we allow some value assignments to assign the same individual to more than one variable. For that matter, there is nothing wrong with a value assignment that assigns the same individual from A to every variable.

The definition of truth relative to a model will now be given in two stages. First, we will give a recursive definition of *true formula of L_1 with respect to a model M and value assignment g*. Then, on the basis of this intermediate definition, we can very simply state the final definition of *true sentence of L_1 with respect to a model M*. (Note that the value assignment g is not to be considered part of the model M – it has nothing to do with how we interpret the *constant* basic expressions of the language.)

The reader may ask at this point how we decide which value assignment to pick. The answer is, it doesn't matter at all, so long as we pick a particular one. The ultimate definition of truth with respect to a model will turn out not to depend at all on which assignment g was initially picked to "compute", as it were, the intermediate truth definition, since the semantic rules systematically make reference to *other* value assignments differing from the original g in specified ways. The formal definitions are as follows:

2. *Semantics of L_1*

A *model for L_1* is an ordered pair $\langle A, F\rangle$ such that A is a non-empty set and F is a function assigning a semantic value to each non-logical constant of L_1

(i.e., names, one-place predicates, and two-place predicates). The set of possible semantic values for names is A; the set of possible semantic values for one-place predicates is $\{1, 0\}^A$; the set of possible semantic values for two-place predicates is $(\{1,0\}^A)^A$, (Here we again use the standard set-theoretic notation "X^Y" to stand for the set of all functions from Y to X.) A *value assignment* g is any function assigning a member of A to each variable of L_1. We abbreviate "the semantic value of α with respect to M and g" as $[\![\alpha]\!]^{M,g}$.

A. Semantic values of basic expressions:
  1. If u is an individual variable of L_1, then $[\![u]\!]^{M,g} = g(u)$.
  2. If α is a non-logical constant of L_1, then $[\![\alpha]\!]^{M,g} = F(\alpha)$.

B. Truth conditions for formulas of L_1 relative to M and g:
  1. If δ is a one-place predicate and α is a term, then $[\![\delta(\alpha)]\!]^{M,g} = [\![\delta]\!]^{M,g}([\![\alpha]\!]^{M,g})$.
  2. If γ is a two-place predicate and α and β are terms, then $[\![\gamma(\alpha,\beta)]\!]^{M,g} = [[\![\gamma]\!]^{M,g}([\![\beta]\!]^{M,g})]([\![\alpha]\!]^{M,g})$.
  3.-7. If ϕ is a formula, then $[\![\neg\phi]\!]^{M,g} = 1$ iff $[\![\phi]\!]^{M,g} = 0$; otherwise, $[\![\neg\phi]\!]^{M,g} = 0$. Similarly for $[\phi \wedge \psi]$, $[\phi \vee \psi]$, $[\phi \rightarrow \psi]$, and $[\phi \leftrightarrow \psi]$.
  8. If ϕ is a formula and u is a variable, then $[\![\forall u\phi]\!]^{M,g} = 1$ iff for every value assignment g' such that g' is exactly like g except possibly for the individual assigned to u by g', $[\![\phi]\!]^{M,g'} = 1$.
  9. If ϕ is a formula and u is a variable, then $[\![\exists u\phi]\!]^{M,g} = 1$ iff for some value assignment g' such that g' is exactly like g except possibly for the individual assigned to u by g', $[\![\phi]\!]^{M,g'} = 1$.

As should be clear, the semantic value $[\![\alpha]\!]^{M,g}$ of any expression α can depend on the particular assignment g only with regard to what values g assigns to variables that are free in α. That is, if $g(u) = g'(u)$ for all variables u that are free in α, then $[\![\alpha]\!]^{M,g} = [\![\alpha]\!]^{M,g'}$. If α is a basic expression, this follows directly from A. When α is a formula, it follows from two facts:

(i) that the free variables of the formulas treated in clauses B.1-7 are exactly those that are free in one or more of the next smaller parts; and

(ii) that the free variables of the formulas treated in clauses B.8-9 are all those except u which are free in the next smaller part and, moreover, the semantic value of the larger formula in B.8-9 is independent of what g assigns to u.

C. We adopt the following truth definition for formulas of L_1 relative to M:

1. For any formula ϕ of L_1, $[\![\phi]\!]^M = 1$ if $[\![\phi]\!]^{M,g} = 1$ for all value assignments g.

2. For any formula ϕ of L_1, $[\![\phi]\!]^M = 0$ if $[\![\phi]\!]^{M,g} = 0$ for all value assignments g.

If a formula ϕ has one or more free variables then it may well be true with respect to some assignments and false with respect to others. In this case its truth or falsity with respect to M is left undefined by C. This last possibility is of no great consequence since it is really only for the *sentences* of L_1 that we are interested in knowing truth values independently of an assignment.[1]

To understand more clearly the workings of clauses B.8 and B.9, we will consider a few simple examples.

For the sake of brevity, we will choose a model for L_1 with a very small domain, *viz.*, the set $\{a, b, c\}$. We choose a set whose members are letters of the alphabet to facilitate the explanation that follows. It is of course somewhat odd to think of a language which can be used to talk about nothing but letters of the alphabet, so the reader may wish to think of this set as consisting of persons or objects of some other kind. However, as before, it is important to keep in mind that the things in the domain A are the *objects* of discourse themselves and not merely some auxiliary *names* of objects.

The *model* M will be the pair $\langle A, F\rangle$, where A is the set $\{a, b, c\}$ and F is as follows:

(3-1)
$$\begin{aligned} F(j) &= a \\ F(d) &= b \\ F(n) &= c \\ F(m) &= a \\ F(M) &= \{a, b, c\} \\ F(B) &= \{b, c\} \\ F(K) &= \{\langle a, a\rangle, \langle a, b\rangle, \langle b, c\rangle\} \\ F(L) &= \{\langle a, c\rangle, \langle b, a\rangle, \langle c, a\rangle, \langle c, c\rangle\}. \end{aligned}$$

(Note that we have specified the semantic values of M, B, K, and L in the form of sets rather than as the corresponding functions.)

Suppose we pick as our initial value assignment g some function that assigns the object c to the variable x, assigns b to y, and assigns a to

z. We will not worry about what g assigns to the infinitely many other variables of L_1 since we will only be concerned with examples containing these three variables. Accordingly, we may represent (the initial part of) g as follows:

$$(3\text{-}2)\quad \begin{bmatrix} x \rightarrow c \\ y \rightarrow b \\ z \rightarrow a \\ \cdots \\ \cdots \\ \cdots \end{bmatrix}$$

Having now given semantic values for all the basic expressions of L_1, we can consider how the truth or falsity of a formula, say $\forall x M(x)$, will be determined with respect to M and g (and ultimately with respect to M alone) by the semantic rules of L_1. The formula in question is built up first by forming $M(x)$ by syntactic rule B.1 and then forming $\forall x M(x)$ by syntactic rule B.8. The semantic rule B.1 tells us that in this case $[\![M(x)]\!]^{M,g} = 1$ (i.e., $M(x)$ is true with respect to M and g) since $g(x) = c$ and $[\![M]\!]^{M,g} = \{a, b, c\}$ (more precisely, $[\![M]\!]^{M,g}$ is the characteristic function of the set $\{a, b, c\}$.) To determine the semantic value of $\forall x M(x)$, we must determine the semantic value of $M(x)$ not just with respect to g but also with respect to all value assignments *like g except for the value assigned to x.* Let us now introduce the following notational convention:

> Notational Convention 4: We use "g_u^e" to indicate the value assignment exactly like g except that it assigns the individual e to the variable u.

Thus in addition to g, we will have to consider the truth or falsity of $M(x)$ with respect to M, g_x^a and with respect to M, g_x^b. Since there are only three individuals in A, there can be only these three distinct variable assignments differing at most in the value assigned to x (g_x^c is identical to g itself). These other two are as follows (where the ellipsis represents exactly the *same* completions of g_x^a and g_x^b as in g.):

(3-3) $g_x^a: \begin{bmatrix} x \to a \\ y \to b \\ z \to a \\ \cdots \\ \cdots \\ \cdots \end{bmatrix}$ $g_x^b: \begin{bmatrix} x \to b \\ y \to b \\ z \to a \\ \cdots \\ \cdots \\ \cdots \end{bmatrix}$

Now we see that $M(x)$ is true with respect to M and any of these three assignments (since a, b, and c are all in $F(M)$), hence $\forall x M(x)$ is true with respect to M, g according to semantic rule B.8. It should also be clear that this sentence would have come out true with respect to any variable assignment g, no matter *which* one we had picked at the outset, since we systematically considered all other assignments assigning different values to x which is the only variable appearing in the formula. Thus in accordance with clause C.1, we drop the reference to g and say simply that $\forall x M(x)$ is *true relative to M.*

Note also that for an existential formula, say $\exists x B(x)$, it likewise does not matter what g we pick originally. For the truth of this formula we require only that *some* value assignment like g (except possibly for the value assigned to x) make $B(x)$ true. It does not matter whether the g we initially pick happens to be one that makes $B(x)$ true or not, as long as there is one that does. Since all bound variables in all formulas will be interpreted through semantic rules B.8 or B.9 sooner or later, the initial choice of g turns out to be irrelevant for all formulas containing only bound variables, hence the possibility of using the definition in the semantic rules in C.

As we mentioned earlier, vacuous quantification will have no semantic consequences. Suppose we have a formula $\forall x M(x)$ and we vacuously quantify it by syntactic rule B.9, giving $\exists y \forall x M(x)$. By the semantic rule B.9, this will be true w.r.t. M, g just in case $\forall x M(x)$ is true w.r.t. M and to some assignment just like g except for the value assigned to y. Now y does not occur in $\forall x M(x)$, so what value y takes on can have no effect on the truth of it; rather, the truth value of $\exists y \forall x M(x)$ with respect to M and this last series of assignments will in every case be the same as the truth value of $\forall x M(x)$ with respect to M and g.

PROBLEM (3-1) Show that vacuous quantification of $B(x)$ by syntactic rule B.8 to yield $\forall y B(x)$ has no effect on the truth value of $B(x)$.

Now consider an example involving two quantifiers: $\forall x \exists y L(x, y)$, which is formed by using syntactic rule B.2, then B.9, then B.8. We note initially that

$L(x, y)$ is false w.r.t. M, g because $\langle c, b\rangle \notin F(L)$. But by semantic rule B.9, $\exists y L(x, y)$ will be true w.r.t. M, g iff we can find some assignment g' differing only in the value assigned to y that makes $L(x, y)$ true w.r.t. M, g'. We might first test g_y^a:

(3-4) $$g_y^a: \begin{bmatrix} x \rightarrow c \\ y \rightarrow a \\ z \rightarrow a \\ \cdots \\ \cdots \\ \cdots \end{bmatrix}$$

Here, $L(x, y)$ turns out to be true (since $\langle c, a\rangle \in F(L)$), so we know without any checking of further values for y that $\exists y L(x, y)$ is true w.r.t. M, g. Now we proceed to the full formula $\forall x \exists y L(x, y)$. By B.8 this formula will be true w.r.t. M, g just in case $\exists y L(x, y)$ is true w.r.t. M and to all g', where g' may differ from g in the value assigned to x. Now we have already checked $\exists y L(x, y)$ for one of the relevant assignments, namely g itself, and found that $\exists y L(x, y)$ is true for it. We then proceed to check g_x^a and g_x^b (exhibited above). To find out whether the formula $\exists y L(x, y)$ is true w.r.t. M, g_x^a, we must in turn ask whether $L(x, y)$ is true w.r.t. some assignment possibly differing from g_x^a in the value assigned to y. (Note that we earlier checked $L(x, y)$ with respect to assignments differing not from g_x^a but from g in the value assigned to y; the difference between the two kinds of assignments is crucial). That is, we will need to check $L(x, y)$ for g_x^a, $[g_x^a]_y^a$, and $[g_x^a]_y^c$ at this stage to see if it is true for at least one of them. These assignments are:

(3-5) $$g_x^a: \begin{bmatrix} x \rightarrow a \\ y \rightarrow b \\ z \rightarrow a \\ \cdots \\ \cdots \\ \cdots \end{bmatrix} \quad [g_x^a]_y^a: \begin{bmatrix} x \rightarrow a \\ y \rightarrow a \\ z \rightarrow a \\ \cdots \\ \cdots \\ \cdots \end{bmatrix} \quad [g_x^a]_y^c: \begin{bmatrix} x \rightarrow a \\ y \rightarrow c \\ z \rightarrow a \\ \cdots \\ \cdots \\ \cdots \end{bmatrix}$$

For the first two of these, $L(x, y)$ comes out false (since neither $\langle a, b\rangle$ nor $\langle a, a\rangle$ is in $F(L)$), but it comes out true for the third (since $\langle a, c\rangle \in F(L)$), so $\exists y L(x, y)$ is true w.r.t. M, g_x^a. Finally, we check the truth of $\exists y L(x, y)$ w.r.t. M, g_x^b, $[g_x^b]_y^a$, or $[g_x^b]_y^c$:

(3-6) $g_x^b: \begin{bmatrix} x \to b \\ y \to b \\ z \to a \\ \cdots \\ \cdots \\ \cdots \end{bmatrix}$ $[g_x^b]_y^a: \begin{bmatrix} x \to b \\ y \to a \\ z \to a \\ \cdots \\ \cdots \\ \cdots \end{bmatrix}$ $[g_x^b]_y^c: \begin{bmatrix} x \to b \\ y \to c \\ z \to a \\ \cdots \\ \cdots \\ \cdots \end{bmatrix}$

The formula $L(x, y)$ is false with respect to the first (because $\langle b, b\rangle \notin F(L)$) but true for the second (because $\langle b, a\rangle \in F(L)$), so we need not bother with the third: we already know that $\exists y L(x, y)$ is true w.r.t. M, g_x^b.

Since we have now investigated $\exists y L(x, y)$ for all possible g' differing from g in the value assigned to x and found it true in these cases, we know by B.8 that $\forall x \exists y L(x, y)$ is true w.r.t. M, g. And since there are no free variables in the formula, we can be assured that it would come out true with respect to M and to any g, hence it is true with respect to M.

As a last example consider the formula $\exists y \forall x L(x, y)$ which is like the previous example but with the scope of the quantifiers reversed. We begin by noting that $L(x, y)$ is false w.r.t. M, g because $\langle c, b\rangle \notin F(L)$. Thus we know already at this point that $\forall x L(x, y)$ cannot be true w.r.t. M, g since there is at least one value assignment (namely, g itself) for which $L(x, y)$ is false. We next move to the question whether $\exists y \forall x L(x, y)$ is true w.r.t. M, g. By the semantic rule B.9, this will be the case iff there is some variable assignment like g except possibly for the value given to y for which $\forall x L(x, y)$ is true. We already know this desired assignment cannot be g itself, so we try others. Is $\forall x L(x, y)$ true w.r.t. g_y^a? If so, then $L(x, y)$ must be true for all assignments like g_y^a except for the individual assigned to x. $L(x, y)$ is true w.r.t. M, g_y^a because $\langle c, a\rangle \in F(L)$, but $L(x, y)$ is false w.r.t. $M, [g_y^a]_x^a$ because $\langle a, a\rangle \notin F(L)$. Therefore, $\forall x L(x, y)$ is false w.r.t. M, g_y^a, and we move on to yet another assignment. Is $\forall x L(x, y)$ true w.r.t. M, g_y^c? $L(x, y)$ is true w.r.t. M, g_y^c, because $\langle c, c\rangle \in F(L)$. Also, $L(x, y)$ is true w.r.t. $M, [g_y^c]_x^a$, because $\langle a, c\rangle \in F(L)$. However, $L(x, y)$ is false w.r.t. $M, [g_y^c]_x^b$ because $\langle b, c\rangle \notin F(L)$. Now we have exhausted all assignments differing from g in the value assigned to y (since g_y^b is g itself) and found $\forall x L(x, y)$ true for none of them. Therefore $\exists y \forall x L(x, y)$ is false w.r.t. M, g. Once again, we can see that this result depended in no way on the g chosen at the outset, so $\exists y \forall x L(x, y)$ is false w.r.t. M by C.2.

If it has not yet become intuitively clear how this procedure works and how it extends to more complex formulas with more quantifiers (or with complex formulas involving sentential connectives and quantifiers), then the reader is encouraged to form further examples and mechanically determine their truth or falsity with respect to M or other constructed finite models by the semantic rules for L_1. Fortunately, it is never necessary to carry out tedious computations of this sort in working with Montague's treatment of English – one's intuitive understanding of formulas in predicate logic is generally sufficient for seeing the point of the English examples. Nevertheless, it is important to keep in mind that a rigorous model-theoretic quantification theory underlies all the quantified formulas to be discussed throughout the rest of the book. This is not the place, however, to discuss mathematical proofs about quantification theory (the so-called *metatheorems* of first-order logic); for these the reader is referred to either Tarski's original treatment (Tarski 1935) or, what is perhaps preferable, to any logic textbook that treats formal semantics of first-order logic (e.g. Van Fraassen 1971, Church 1956, Quine 1951).

As a final comment to our presentation of L_1 we note that the definitions given in the preceding chapter of validity, entailment, etc. can be carried over directly to this new language. Since the dependence on a value assignment has been gotten rid of when we consider the truth values of sentences, we can continue to say that a valid sentence is one which is true with respect to every model, a contradictory sentence is false with respect to every model, and so on.

PROBLEM (3-2) Show by a detailed consideration of the relevant value assignments that the sentence $\forall x \forall y [L(x, y) \rightarrow L(y, x)]$ is false with respect to the given model M. Find a model $M' = \langle A, F' \rangle$ such that the sentence is true with respect to M'.

PROBLEM (3-3) Reformulate the syntactic and semantic rules of L_1 so that the logical connectives are assigned to basic categories rather than being introduced syncategorematically. What difficulties arise in attempting to make the same move with "$\forall$" and "$\exists$"?

PROBLEM (3-4) Show that the sentences $\forall x [B(x) \leftrightarrow B(x)]$ and $\exists x [M(x) \rightarrow \forall y M(y)]$ are valid.

II. THE LANGUAGE L_{1E}

Now that we have examined a logical language which allows quantification over individual variables, we next want to consider how our English-like fragment, L_{0E}, from the preceding chapter might be enlarged to accommodate similar syntactic and semantic processes. Languages L_0 and L_{0E} were virtually isomorphic syntactically, and they were deliberately chosen in this way to provide simple illustrations of the application of model-theoretic semantics to both natural and formal languages. When we consider quantification in natural language as opposed to a formal language such as L_1, however, we see that the two are different in important respects. Anyone who has taken a course in symbolic logic knows that some degree of skill is required to "translate" English sentences into first-order predicate logic. For example, an English sentence such as "Every man walks" is to be rendered in predicate logic as something like $\forall x [M(x) \rightarrow W(x)]$, a formula which contains a logical connective and three instances of a variable having no direct counterparts in the English sentence. Further, both the common count noun "man" and the intransitive verb "walks" have been represented in the logical formula by one-place predicates. If we then go on to reflect on the fact that there are many English quantifiers such as "most," "few," "many," and "much" which have no ready correspondents in predicate logic at all, it becomes clear that the syntax and semantics of quantification in English (or indeed of any other natural language) cannot be any simple isomorphism of L_1.

Quantificational phenomena in English (and in all other natural languages) are in fact so complex that there are still many problems in this area which have not been solved. Indeed, the importance of Montague's paper 'The Proper Treatment of Quantification in Ordinary English' lies in part in the fact that it represents an important advance in the direction of solving some of these problems. In the fragment that we will next construct, we will attempt to keep matters as simple as possible. For example, we will consider only the quantifiers "every," "some," and "the," and we will avoid entirely the problems raised by mass nouns, plural count nouns, relative clauses, and a host of other constructions. This fragment, however, will serve as an indication of one sort of approach to quantification in natural language, and in fact it is quite similar to the framework which Montague adopted in a paper which antedates PTQ, viz., 'English as a Formal Language.'

If we compare simple English quantificational statements with their translations into predicate logic, we note some obvious differences:

(3-7) a. Every student walks $\forall x [S(x) \rightarrow W(x)]$
b. Some student walks $\exists x [S(x) \wedge W(x)]$

While quantification in English is expressed by a determiner combined with a common noun to form a noun phrase, the effect in predicate logic is achieved by two syntactically independent devices: *variables*, which play the same syntactic role as names, and *quantification rules* which are later used to form a new formula from a formula. In the English examples above, however, nothing corresponding to a variable is evident. Nevertheless, there are at least some noun phrases in English that do seem to function as variables do in logic, namely certain pronouns, such as the underlined pronouns in (3-8):

(3-8) a. Every Englishman loves himself
b. Every Englishman loves his mother
c. Every Englishman believes that he is honorable

It has often been observed that from a semantic point of view at least, these pronouns do not merely serve as syntactic substitutes for their antecedents, for the sentences in (3-9) are not synonymous with their counterparts in (3-8):

(3-9) a. Every Englishman loves every Englishman
b. Every Englishman loves every Englishman's mother
c. Every Englishman thinks that every Englishman is honorable

Rather, we can paraphrase the meaning of (3-8a) correctly if we say that sentence (3-10a) below is true for every value of v_1 that is an Englishman (and similarly for (3-10b) and (3-10c)):

(3-10) a. v_1 loves v_1
b. v_1 loves v_1's mother
c. v_1 thinks that v_1 is honorable

Since this kind of paraphrase is obviously reminiscent of our semantic clause for $\forall x\phi$, these examples suggest the possibility that the sentences in (3-8) could be produced in an English-like formal language by using the quasi-sentences in (3-10) as an intermediate step, letting v_1 play the semantic role of a variable and letting the semantic rule corresponding to the syntactic conversion of (3-10) into (3-8) work like the semantic rule for $\forall x\phi$ in L_1. This syntactic "quantification" process that we will introduce into L_{1E} will have to do two things: the leftmost occurrence of the variable v_1 (or other variable) will have to be replaced by a full noun phrase such as *Every Englishman*, while the subsequent occurrence of v_1 in the sentence must be turned into a pronoun, such as *himself, his*, or *he*.

But now what of the simpler examples in (3-7)? Though we might treat them in a syntactically simpler way (and in fact, we will eventually see this possibility realized in the PTQ English syntax), note that this two-step process needed for (3-8) would work just as well for these. That is, we can produce *Every fish snores* from v_1 *snores* by the same "replacement" operation (corresponding semantically to variable binding) as suggested above, the only difference here being that no second occurrence of the variable v_1 need be involved. Thus we will formulate our "quantification" rule for L_{1E} as replacing the first occurrence of a variable by a noun phrase containing a determiner *every* (or *some* or *the*) and at the same time replacing all subsequent occurrences of that variable, *if there are any*, with a pronoun (or, as we shall choose to do for simplicity in L_{1E}, a "pronoun substitute").

We specify the syntax of our English fragment L_{1E} as a set of formation rules rather than in the form of a phrase-structure grammar. The reader will also note that the syntax of our new English-like fragment contains "variables" v_1, v_2, etc. as basic expressions of category N. These are of course not really basic expressions of English, but they are essential to the approach to quantification we will assume here, and it will do no harm to include them so long as we distinguish, as we did in the case of L_1, between formulas and sentences. Thus, we will define the auxiliary concept of an "English formula" and then define in terms of this the notion of "English sentence."

In contrast to the distinction between formulas and sentences of L_1, however, English sentences will contain no occurrences of variables at all, while they may occur in English formulas. This reflects the different way in which variables will be treated in the English-like fragment – variables do not become *bound* in the way that the occurrence of x, for example, becomes bound when $B(x)$ is quantified by $\forall x$ or $\exists x$ in L_1; rather, the variables of L_{1E} will be *replaced* by lexical items of English. For example, by formation rule B5 below, we can form from the English formula v_1 *snores* and the common noun *man* the English sentence *Every man snores*. Similarly, from the formula v_1 *snores or* v_1 *is-boring* and the common noun *man* the same rule licenses the formation of the sentence *Every man snores or that man is-boring*.

Here is the formal specification of the syntax of L_{1E}:

1. *Syntax of* L_{1E}

A. The basic expressions of L_{1E} are as follows:

1. *Hank, Liz, Sadie* are constants of category N.
2. $v_1, v_2, v_3, \ldots$ are variables of category N.
3. *sleeps, snores, is-boring* are constants of category V_i.
4. *loves, hates, is-taller-than* are constants of category V_t.
5. *man, woman, fish* are constants of category CN (common noun).

6. *it-is-not-the-case-that* is a (logical) constant of category *Neg*.
7. *and*, *or* are (logical) constants of category *Conj*.

B. The formation rules of L_{1E} are as follows:
1. If α is a V_t and β is an N, then $\alpha\beta$ is a V_i.
2. If α is a V_i and β is an N, then $\beta\alpha$ is a *For* (formula).
3. If α is a *Neg* and ϕ is a *For*, then $\alpha\phi$ is a *For*.
4. If α is a *Conj* and ϕ and ψ are *For*'s then $\phi\alpha\psi$ is a *For*.
5. If α is a *CN*, u is a variable, and ϕ is a *For* containing at least one occurrence of u, then ϕ' is a *For*, where ϕ' comes from ϕ by replacing the left-most occurrence of u by *every* α and each subsequent occurrence of u by *that* α.
6. If α is a *CN*, u is a variable, and ϕ is a *For* containing at least one occurrence of u, then ϕ' is a *For*, where ϕ' comes from ϕ by replacing the left-most occurrence of u by *some* α and each subsequent occurrence of u by *that* α.
7. If α is a *CN*, u is a variable, and ϕ is a *For* containing at least one occurrence of u, then ϕ' is a *For*, where ϕ' comes from ϕ by replacing the left-most occurrence of u by *the* α and each subsequent occurrence of u by *that* α.

C. If ϕ is a *For* by the rules in *A* and *B* above and contains no instances of any variable, then ϕ is an *S* (sentence).

Rules B5, 6, and 7, will, for example, allow us to form from the common noun *fish* and the formula *Hank loves* v_6 the sentences *Hank loves every fish*, *Hank loves some fish*, and *Hank loves the fish*, respectively. Note that the correct interpretation of these rules requires careful attention to the object language/meta-language distinction. In B5, for example, we are to replace the left-most occurrence of the variable not by the sequence "*every* α" to produce, say, "*Hank loves every* α"; rather, what is substituted for the variable occurrence is the sequence consisting of the word *every* followed by whatever word it is in the object language that is taken as the value of α.

Note how different the derivations are in L_{1E} of *Hank snores* and *Every man snores*. The former is derived in the way indicated by the following analysis tree:

The derivation of the other sentence may be displayed in the following fashion:

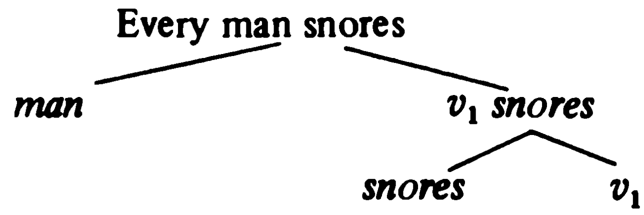

The grammar of L_{1E} has no category of noun phrases containing both *Hank* and *every man*. Indeed *every man* is not a phrase at all in L_{1E}! This is obviously a defect in our syntactic analysis of the language. Our reason for writing the grammar this way will become obvious when we state the semantic rules corresponding to syntactic rules B1-7.

Our syntax here follows Montague in using *that* plus another instance of the common noun to replace all but the first occurrence of a variable. This is a compromise which avoids the problems connected with personal vs. reflexive pronouns and with gender agreement. Thus, our fragment generates from *man* and v_1 *loves* v_1 the rather unnatural sentence *Every man loves that man* instead of *Every man loves himself*, and similarly from *man* and *Hank loves* v_1 *or Liz loves* v_1 we derive *Hank loves some man or Liz loves that man* rather than *Hank loves some man or Liz loves him*. In this connection, we should also point out another deficiency in our fragment vis-a-vis ordinary English. By our syntax, the sentence *Every man loves Hank and Liz loves that man* can only arise from a formula such as v_1 *loves Hank and Liz loves* v_1 (by rule B5). There is no possibility that the phrase *that man* could arise, as it were, as an epithet for Hank. Indeed, in our fragment we cannot generate sentences such as *Hank snores and that man is-boring* at all.

A comment is also in order concerning one peculiar formal property of our syntax: every sentence containing *every*, *some*, or *the* will have infinitely many distinct syntactic analyses. *Every man sleeps*, for example, could have been formed by rule B5 out of *man* and any one of the formulas v_1 *sleeps*, v_2 *sleeps*, v_3 *sleeps*, etc., etc. This may seem worrisome if one is accustomed to thinking that distinct syntactic derivations must receive distinct semantic interpretations, but as we shall see, this is not so in the semantics we give for L_{1E}. The sentence *Every man sleeps* will not turn out to be infinitely many ways ambiguous; it will turn out not to be even two ways ambiguous semantically. Thus, this proliferation of syntactic derivations need not bother us if we are willing to accept it as a harmless and trivial formal property of our system.

There will, however, arise sentences which will have genuinely nonequivalent syntactic derivations and to which we will want to assign semantic values in more than one way. For example, the formula v_1 *loves* v_2 can give rise to the sentence *Some man loves every woman* in two essentially different ways. First, by rule B5 we could form v_1 *loves every woman* and then by rule B6 form *Some man loves every woman*, or else we could apply these rules in the opposite order, forming first *Some man loves* v_2 and then the sentence in question. As we will see, the semantics will specify truth conditions for this sentence in two distinct ways corresponding to the two derivations just mentioned. The truth conditions will state that for the sentence derived in the first way it is true just in case there is someone who is a man and such that he loves every woman. The truth conditions for the sentence when derived in the second way will correspond to the other "reading" in which for each woman there is some man or other, not necessarily the same man for all women, who loves her. Thus, our fragment provides the means for capturing, at least for some sentences, differences in meaning that depend on the relative scope of quantifiers. We note also that our fragment allows nonequivalent derivations for sentences involving at least two connectives, just as the syntax of L_{0E} did; e.g., *Sadie sleeps or Liz is-boring and Hank snores.*

Thus in L_{1E} as in L_{0E} we will need to assign semantic values to analysis trees of the kind illustrated above. These structures, but not the expressions of which the trees are structural descriptions, contain enough information about syntactic derivations to permit the semantic rules to retrace them.

2. *Semantics of* L_{1E}

It is clear from the syntactic form of sentences containing quantifiers in L_{1E} that the truth conditions will have to be given in a form different from those for the quantifiers $\forall$ and $\exists$ in L_1. For example, in determining the truth value of a sentence such as *Every man sleeps* it will not do to say that it is true just in case the formula v_1 *sleeps* is true for all assignments of values to the variable v_1. That would give us the truth conditions for something like *Everyone sleeps* (i.e., every individual in the universe of discourse sleeps). What we want, rather, is something like "v_1 *sleeps* is true for every assignment of an individual to the variable v_1 such that that individual is a man." The semantic rules below thus treat quantification as "restricted" to a subdomain which is indicated by the common noun with which the quantifier is associated. This, in turn, indicates the sort of semantic values we will want to assign to common nouns, *viz.*, a function from individuals to truth values. This characteristic function of some set will thus represent the set of men, women, fish or whatever, just as the same sort of characteristic function indicates the set of snorers, sleepers, etc.

Here is the formal specification of the semantics of L_{1E}:

A model for L_{1E} is an ordered pair $\langle A, F\rangle$ such that A is a non-empty set and F is a function assigning a semantic value to each non-logical constant and which assigns to *it-is-not-the-case-that*, *and*, and *or* the semantic values appropriate to the logical connectives $\neg$, $\wedge$, and $\vee$ respectively. The domains of possible semantic values for expressions of each syntactic category are given in the following table:

Category	*Set of Possible Semantic Values*
N	A
For	$\{1, 0\}$
V_i	$\{1, 0\}^A$
CN	$\{1, 0\}^A$
V_t	$(\{1, 0\}^A)^A$

A value assignment g is a function assigning to each variable a semantic value of the appropriate sort (here, as in L_1, assigning an individual to each variable).

A. Semantic values of basic expressions:
1. If u is a variable, then $[\![u]\!]^{M,g} = g(u)$.
2. If α is a non-logical constant, then $[\![\alpha]\!]^{M,g} = F(\alpha)$.
3. If α is a logical constant (member of *Neg* or *Conj*) then $[\![\alpha]\!]^{M,g}$ is as follows:

$$[\![\textit{It-is-not-the-case-that}]\!]^{M,g} = \begin{bmatrix} 1 \to 0 \\ 0 \to 1 \end{bmatrix}$$

$$[\![\textit{and}]\!]^{M,g} = \begin{bmatrix} 1 \to & \begin{bmatrix} 1 \to 1 \\ 0 \to 0 \end{bmatrix} \\ 0 \to & \begin{bmatrix} 1 \to 0 \\ 0 \to 0 \end{bmatrix} \end{bmatrix} \qquad [\![\textit{or}]\!]^{M,g} = \begin{bmatrix} 1 \to & \begin{bmatrix} 1 \to 1 \\ 0 \to 1 \end{bmatrix} \\ 0 \to & \begin{bmatrix} 1 \to 1 \\ 0 \to 0 \end{bmatrix} \end{bmatrix}$$

B. Truth conditions of formulas of L_{1E} relative to M and g:
1. If α is a V_t and β is an N, then $[\![\alpha\beta]\!]^{M,g} = [\![\alpha]\!]^{M,g}([\![\beta]\!]^{M,g})$.
2. If α is a V_i and β is an N, then $[\![\beta\alpha]\!]^{M,g} = [\![\alpha]\!]^{M,g}([\![\beta]\!]^{M,g})$.
3. If α is a *Neg* and ϕ is a *For*, then $[\![\alpha\phi]\!]^{M,g} = [\![\alpha]\!]^{M,g}([\![\phi]\!]^{M,g})$.
4. If α is a *Conj* and ϕ and ψ are *For*'s, then $[\![\phi\alpha\psi]\!]^{M,g} = [[\![\alpha]\!]^{M,g}([\![\phi]\!]^{M,g})]([\![\psi]\!]^{M,g})$.
5. If α is a *CN*, u is a variable, and ϕ is a *For* containing at least one

occurrence of u, then for ϕ' as in syntactic rule B5, $[\![\phi']\!]^{M,g} = 1$ iff for all value assignments g_u^e such that $[\![\alpha]\!]^{M,g}(e) = 1$, $[\![\phi]\!]^{M,g_u^e} = 1$.

6. If α is a *CN*, u is a variable, and ϕ is a *For* containing at least one occurrence of u, then for ϕ' as in syntactic rule B6, $[\![\phi']\!]^{M,g} = 1$ iff for some value assignment g_u^e such that $[\![\alpha]\!]^{M,g}(e) = 1$, $[\![\phi]\!]^{M,g_u^e} = 1$.
7. If α is a *CN*, u is a variable, and ϕ is a *For* containing at least one occurrence of u, then for ϕ' as in syntactic rule B7, $[\![\phi']\!]^{M,g} = 1$ iff there is exactly one e in A such that $[\![\alpha]\!]^{M,g}(e) = 1$, and furthermore $[\![\phi]\!]^{M,g_u^e} = 1$.

As already noted, we need really to assign semantic values to analysis trees rather than directly to strings of basic expressions. The reader should understand the rules of A and B in that way. It is only to avoid the distraction of fussy details that we have formulated them in an oversimplified fashion. The same remark applies to C.

C. Truth conditions for sentences of L_{1E} relative to M:
1. For any sentence ϕ of L_{1E}, $[\![\phi]\!]^{M} = 1$ if for all value assignments g, $[\![\phi]\!]^{M,g} = 1$.
2. For any sentence ϕ of L_{1E}, $[\![\phi]\!]^{M} = 0$ if for all value assignments g $[\![\phi]\!]^{M,g} = 0$.

Let us now illustrate the workings of these rules by choosing a particular model and determining the semantic values of some sentences with respect to that model. We choose A as for L_{0E} in the preceding chapter, i.e., A = {Anwar Sadat, Queen Elizabeth II, Henry Kissinger}. Further, we let F assign to all the non-logical constants of L_{1E} which are in L_{0E} just the same values that we chose for those constants in model M_0 of the preceding chapter, i.e., F(*Hank*) = Henry Kissinger, etc. (pp. 26–33). To complete the assignment of values to the constants of L_{1E} we let the values of *man*, *woman*, and *fish* be as follows:

(3-8) $[\![man]\!]^{M,g} =$

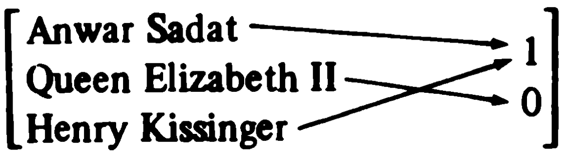

(3-9) $[\![woman]\!]^{M,g} =$

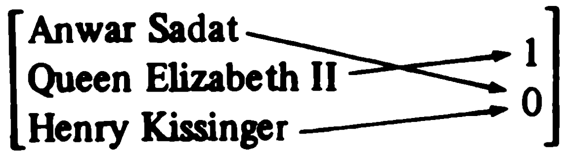

(3-10) $[\![fish]\!]^{M,g} =$

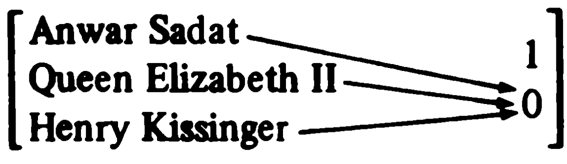

Let us also take as the value assignment g the following function:

$$(3\text{-}11) \quad g = \begin{bmatrix} v_1 \rightarrow \text{Anwar Sadat} \\ v_2 \rightarrow \text{Henry Kissinger} \\ v_3 \rightarrow \text{Queen Elizabeth} \\ \cdots\cdots\cdots\cdots \\ \cdots\cdots\cdots\cdots \end{bmatrix}$$

where the dots indicate some completion of the function whose exact nature we need not be concerned with. For concreteness, let us suppose that all other variables are assigned the value Queen Elizabeth II.

There is nothing new in sentences of L_{1E} that belong to L_{0E} also: their semantic values will be just as before. Therefore, let us consider straight off the semantic value of the sentence *Every man snores* with respect to this model and value assignment. Recall that syntactically this sentence is generated by first combining a variable, say v_2, with the intransitive verb *snores* to give the formula v_2 *snores*. Then by syntactic rule B5 the common noun *man*, the variable v_2, and the formula v_2 *snores* combine to yield *Every man snores*. Our semantic computations will then parallel these steps. First, semantic rule B2 says that $[\![v_2 \textit{ snores}]\!]^{M,g} = [\![\textit{snores}]\!]^{M,g}([\![v_2]\!]^{M,g})$; that is, the function given in (2-13) applied at the argument Henry Kissinger, the latter being $g(v_2)$. Thus, $[\![v_2 \textit{ snores}]\!]^{M,g} = 0$. Next we determine the semantic value of *Every man snores* (with respect to M and g) by semantic rule B5 on the basis of the semantic values of *man*, v_2, and v_2 *snores*. The semantic value of *man* with respect to M and any g' is given in (3-8) above. To apply semantic rule B5 we need first to find all e such that $[\![\textit{man}]\!]^{M,g}(e) = 1$, that is, all the individuals in A who are men according to this model and value assignment. By (3-8) we see that this is just Anwar Sadat and Henry Kissinger. Thus, we are to construct all value assignments which are like g except that the values assigned to v_2 are to range over the set {Anwar Sadat, Henry Kissinger}. Since g as it happens to have been chosen already assigns v_2 the value Henry Kissinger, we have only one other assignment to consider, namely:

$$(3\text{-}12) \quad g_{v_2}^{\text{Anwar Sadat}} = \begin{bmatrix} v_1 \rightarrow \text{Anwar Sadat} \\ v_2 \rightarrow \text{Anwar Sadat} \\ v_3 \rightarrow \text{Queen Elizabeth II} \\ \cdots\cdots\cdots\cdots \\ \cdots\cdots\cdots\cdots \end{bmatrix}$$

(where the dots indicate the same completion as for *g*). Now the rule tells us that *Every man snores* receives the value 1 iff v_2 *snores* receives the value 1 under *both* the above value assignments. Thus we see that we could have actually stopped after examining the value assignment *g* itself, since it makes v_2 *snores* false, and this would be sufficient for us to conclude that *Every man snores* is false (with respect to this model and value assignment). Since *Every man snores* contains no instances of free variables, it will be false in this model no matter which value assignment we start with. Thus, by rule C2 above, $[\![Every\ man\ snores]\!]^M = 0$.

The reason we distorted English syntactic structure so much in the grammar of L_{1E} was to allow ourselves to formulate semantic rules B.5-7. If we had generated a category of noun phrases including *every man, some man, Hank, etc.*, then we would have found ourselves in a bind when trying to assign a suitable denotation to *some man,* for example. The requisite semantic value cannot be an individual. To see this, note that if it were, then what individual it is would have to depend only on the model and not on a value assignment – since *some man* does not contain any free variables. But then *Some man snores and it-is-not-the-case-that some man snores* (cf. the more colloquial "Some man snores and some doesn't snore") would have to be false in any model – for the same reason that *Hank snores and it-is-not-the-case-that Hank snores* must be. Nor would it help to assign a set of individuals as the semantic value of *some man*. An adequate semantic value will be developed in the next chapter. In the meantime, we have circumvented this particular semantic pitfall by providing a strange syntactic structure – merely for didactic purposes. It would be a good idea for the reader to verify that the problem sentence, *Some man snores and it-is-not-the-case-that some man snores*, is true (on one derivation) in the model under consideration.

In determining the semantic value of a quantified sentence, given a particular model and the value assignment, there is no need of course to go through such excruciating detail in practice as we did above. One can instead run quickly through a chain of reasoning like the following: What is the value of *man* according to the model: Answer: {Anwar Sadat, Henry Kissinger}. Are all the members of this set in the value of *snores*? Answer: No. Therefore, *Every man snores* is false in this model. Similarly, the truth value of *Some man snores* is easily seen to be true by determining that at least one of the members of the set {Anwar Sadat, Henry Kissinger}, namely Anwar Sadat, is in the semantic value of *snores*.

Now consider the example *Every fish snores*. According to the model we have chosen, there are no fish in the universe of discourse (the value of ***fish*** is the characteristic function of the null set). Thus, in applying semantic rule B5, there will be no value assignments g_u^e such that $[\![\textit{fish}]\!]^{M,g}(e) = 1$. Thus, *Every fish snores* will be vacuously true (with respect to this model) since there are no *e* for which $[\![v_2 \textit{ snores}]\!]^{M, g_{v_2}^e}$ must be true. Alternatively, one can reason as follows: Are all the members of the set which is the value of *fish* also in the semantic value of *snores?* Answer: Yes, because there are no members in the former set.

Note, however, that semantic rule B6 requires for the truth of a sentence such as *Some fish snores* that there be some value assignment g_u^e such that $[\![\textit{fish}]\!]^{M,g}(e) = 1$, etc, and since there is no such assignment, *Some fish snores* will receive the value 0 (in this model). This result and the one just mentioned are in accord with one traditional way of construing the meanings of English quantifiers; that is, "Every *X Y*'s" does not entail the existence of any individuals who are *X*, while "Some *X Y*'s" does. It is a consequence of this way of looking at these quantifiers that a sentence of the form "Every *X Y*'s" does not logically imply "Some *X Y*'s" since in case there are no individuals who are *X*, the former will be true but the latter false. In such a case also a sentence of the form "It-is-not-the-case-that every *X Y*'s" will be false. On an alternative view, "Every *X Y*'s" is taken to be false if there are no individuals who are *X*. Then such a sentence would logically imply "Some *X Y*'s", and the sentence "It-is-not-the-case-that every *X Y*'s" would be true. Another standard approach is to postulate a third value which is neither truth nor falsity and say that "Every *X Y*'s" has this third value (or equivalently, has no truth at all) when there is nothing in the semantic value of *X*; in some views one says the same of "It-is-not-the-case-that-every *X Y*'s." A discussion of all the alternatives would entangle us in various intricacies surrounding the notion (or notions) of "presupposition," and that would lead us too far afield at this point. For some references, see Chapter 9.

Similar considerations arise, however, in connection with the treatment of *the* in semantic rule B7. A sentence such as *The man snores* receives the truth value 1 just in case there is exactly one man and that man snores (with respect to the model and value assignment); otherwise, *The man snores* is false. Thus in the model chosen for our illustration, this sentence is in fact false since there are exactly two men. Similarly, *The fish snores* is false inasmuch as there is not, according to the model, exactly one fish who snores, there being no fish at all. We have thus implicitly adopted in semantic rule B7 Russell's theory of definite descriptions (Russell, 1905), according to which a sentence of the form "The *X* *Y*'s" logically implies that there is exactly one individual who is *X* and that individual *Y*'s. Thus, if there is not exactly one individual who is *X*, or if there is, but that individual does not *Y*, then "The *X* *Y*'s" is false. An alternative view – the one adopted for example, in (Strawson, 1950) – is that "The *X* *Y*'s" has no truth value in case there is not exactly one individual who is *X*. In giving the semantic rules as we have, we do not wish to be regarded as espousing a particular point of view on these rather complex issues. We have simply chosen one alternative for the sake of specificity; this alternative happens also to be the one adopted by Montague in PTQ, so the reader will be on a bit more familiar ground when we take up that system in Chapter 7.

PROBLEM (3-5): Determine the truth value of each of the following sentences according to the assumed model for L_{1E}. Which of them, if any, have truth values that differ according to the order in which syntactic rules have been applied?

a. Liz loves every man.
b. The woman hates Hank.
c. Some man sleeps and every woman is-boring.
d. Every man is-taller-than that man.
e. It-is-not-the-case-that some woman hates that woman.

In order to show that a sentence such as *Some man loves every woman* is, according to our system, a case of genuine ambiguity, we will have to choose a slightly more populous model. Let us add to the domain of discourse the individual Jacqueline Onassis, and let us choose the denotations of *man*, *woman* and *loves* in such a way that the men are, as before, just Henry Kissinger and Anwar Sadat, the women are Queen Elizabeth II and Jacqueline Onassis, and the semantic value of *loves* is now given (as shorthand for the appropriate function) as the following set of ordered pairs:

(3-13) $[\![loves]\!]$ = {⟨Henry Kissinger, Jacqueline Onassis⟩, ⟨Anwar Sadat, Queen Elizabeth II⟩}

(Recall that in this representation the first member of the ordered pair is taken as standing in the relation to the second member).

Now let us consider a syntactic derivation of the sentence *Some man loves every woman* which begins with v_1 and the transitive verb *loves*. We assume that the value assignment g is as in (3-11). Thus *loves* v_1 will receive the null set as semantic value (actually the characteristic function of this set) since $g(v_1)$ = Anwar Sadat and in this model no one loves Anwar Sadat. Next, we form v_2 *loves* v_1, and this receives the truth value 0, since $g(v_2)$ = Henry Kissinger and he does not love Anwar Sadat. We are now ready to insert the quantified expressions and show that the resulting sentence is assigned a different truth value depending on the order in which the phrases are inserted.

From *man* v_2, and v_2 *loves* v_1 we form by syntactic rule B6 the formula *Some man loves* v_1. This is assigned the truth value 0 since, according to our model and the value assignment g, there is no man who loves Anwar Sadat. Next, from *woman*, v_1, and *Some man loves* v_1 we form *Some man loves every woman.* To determine the truth value of this, we must find the truth value of *Some man loves* v_1 for every assignment in which the value of v_1 is a woman. When v_1 is assigned the value Queen Elizabeth II, *Some man loves* v_1 is true, since, by semantic rule B6, we determine that v_2 *loves* v_1 is true for at least one value of v_2 which is a man, namely, Anwar Sadat. When v_1 is assigned the value Jacqueline Onassis, *Some man loves* v_1 is true, since v_2 *loves* v_1 is true when v_2 takes on the value Henry Kissinger, who is a man. Thus, the semantic rules in connection with this syntactic derivatin of the sentence *Some man loves every woman* assign the sentence the semantic value 1. Note that the situation that was necessary to ensure this result can be described as "for each woman there is some man or other who loves her."

Now consider the different syntactic derivation in which we first form from *woman* v_1, and v_2 *loves* v_1 the formula v_2 *loves every woman* by syntactic rule B5. This is assigned the value 0 since Henry Kissinger does not love every woman; specifically, he doesn't love Queen Elizabeth II. Now we form *Some man loves every woman* from *man*, v_2, and v_2 *loves every woman* by syntactic rule B6. Here we must let v_2 take on successive values from the set of men and determine whether for any of these individuals it is true that he loves every woman. We already know that it is not true of Henry Kissinger, and when we let the value of v_2 be Anwar Sadat, we see that it is not true of him either (he doesn't love Jacqueline Onassis). Therefore, on this derivation the sentence *Some man loves every woman* is false, and the assumed situation that led to this being so can be expressed as "there is no man such that he loves every woman."

The sentence in question is intuitively ambiguous (at least for most speakers of English) in just the way countenanced by our semantics, so we might justifiably find some satisfaction in this approach to quantification in English. Of course it remains to be seen whether our good fortune will hold up as more complex syntactic constructions are added to the fragment. And one may very well wonder whether we should have to generate the sentence in two syntactically nonequivalent ways, as we did in order to provide a basis for the semantic ambiguity.[2] We will return to some of these questions in connection with our discussion of PTQ in Chapter 7.

A final point. In Chapters 2 and 3 we have repeatedly stressed that model-theoretic semantics proceeds by associating real-world objects with linguistic expressions. Our purpose has been to emphasize the important point that model theory capitalizes on the use of language to talk about things and incorporates the potential for this in a fundamental way. This characteristic of model theory tempts some people to think that the theory embodies too simplistic a notion of the way actual users of language connect expressions with real-world objects. Let us take this opportunity to head off such misapprehensions.

We illustrate one form of the error by formulating an objection someone might make. One might take the position that *every mammal suckles its young* could not possibly mean the truly universal proposition we have analyzed it as meaning, since

(a) human beings can mentally grasp its meaning;
(b) people can, moreover, have adequate grounds for asserting the sentence; but
(c) they may very well not know of the existence of some mammal and then could not be talking about mammals they were not aware of.

The fallacy in the objection lies in the falsity of point (c), which is clearly not a logical consequence of (a) and (b). Human beings certainly can grasp the meaning of universal sentences, and sometimes have adequate grounds for asserting them. But, "adequate" does not mean infallible here anymore than it does elsewhere. If there are mammals a speaker is unaware of, his ignorance of their existence does not excuse him of responsibility in case some of them turn out not to suckle their young – any more than if he were mistaken about the feeding habits of mammals he knew existed. If someone points out mammals he didn't know of which do not suckle their young, he would have to admit he had made a mistake and withdraw his assertion. This shows that this statement applies to all mammals, not only the ones he knew about at the time he asserted it. He could afterwards maintain at best that every mammal in a restricted set suckles its young, and only if the restriction in his statement excluded the nonsuckling mammals he has learned of could he have adequate grounds for this assertion. Thus it is correct to maintain that *every* makes truly universal statements. The mistake our objector made was in thinking that adequate grounds for asserting a sentence suffice to insure that the sentence is true.

Our human imperfections in correctly making the connection of our language to the world do not in any way negate the fact that our understanding of many words and phrases is grounded in exactly that connection. Precisely because this is so, it is appropriate to criticize speakers when they try to make the connections and do so incorrectly. But also because we intend to use language in terms of its real connections, it is possible for us to employ it in an entirely mental way when that is useful. Language can help us think about what it would be like if things were not as they really are. It can help us represent how things might be or might have been. This is important not only in every day activities like planning to make things different, but also in giving scientific explanation of the way things in fact are. Such "detached" uses of language do not reveal an inadequacy of model-theoretic semantics. If the world *were* as it is imagined, expressions *would* connect to objects just as the model has them do. Thus these imaginative uses of language actually illustrate a strength of model-theoretic semantics rather than a weakness.

EXERCISES

1. Reformulate the syntax and semantics of L_1, replacing the basic expression $\neg, \wedge, \vee$, etc. by the "Polish" connectives N, K, A, etc., treated as basic expressions (see Exercise 3, Chapter 2).

2. Show that in syntactic rules B5, 6, and 7 and also in the corresponding semantic rules for L_{1E} the condition that ϕ contain at least one occurrence of the variable u is essential in order to avoid absurd results for sentences such as *Hank snores* which contain no quantifiers.

3. What difficulties arise in attempting to specify the syntax of L_{1E} by means of a context-free phrase-structure grammar? Readers with sufficient background in mathematical linguistics may want to construct a proof that no context-free phrase-structure grammar generating this language is possible.

4. Construct a variant of the syntax and semantics of L_{1E} in which names as well as quantifiers + common ncuns can be substituted for variables. For example, from *Hank* and v_1 *snores* one could form *Hank snores* (This sentence will also be generated in the usual way by syntactic rule B2). Write the syntactic rules in such a way that the phrase *that person* is substituted for all occurrences of a given variable other than the left-most; i.e., *Hank* + v_1 *loves* v_1 yields *Hank loves that person*, and *Hank* + v_1 *snores or Liz loves* v_1 yields *Hank snores or Liz loves that person*. Arrange the semantics in such a way that *Hank snores* has the same truth conditions under either derivation.

5. Show that under the syntax and semantics given for L_{1E} the sentences *Every man loves every man* and *Every man loves that man* have different truth conditions, and further, that the former logically entails the latter.

6. Show that under the syntax and semantics given for L_{1E} the sentence *Every man sleeps or that man snores* does not logically entail the sentence *Every man sleeps or every man snores.*

NOTES

[1] It is a technical consequence of this definition, C., that even certain formulas with free variables will count as true (in the unrelativized sense), namely those in which the formula comes out true no matter what values are assigned to its free variable(s). In other words, a formula with free variables that comes out true by this definition is one that would also be true just in case enough universal quantifiers were prefixed to it to bind all its free variables. Montague in fact exploited this technical consequence in PTQ by omitting universal quantifiers from some of his "meaning postulates": the intent is clearly that these postulates are to be read as if these additional universal quantifiers were present.

In his 'Universal Grammar,' however, Montague used the definition of *true relative to a value assignment g* in a different way. As we will explain in Chapter 8, the function *g* is there thought of as serving the additional purpose of supplying values for deictic pronouns (which are treated as free variables). That is, a formula like *He walks* (i.e. *x walks*) is considered now true, now false, depending on the context in which it is used (i.e. depending on which particular *g* is chosen as "context"). The interpretation of bound variables is unaffected by this new view of *g*. For a detailed analysis of non-anaphoric pronouns in English based on this idea, see Cooper (1979a).

[2] Linguists who consider syntax autonomous find it highly suspect to force syntactic rules to produce ambiguities that are motivated solely by semantic considerations. This particular problem about L_{1E} is not a result of the already noted syntactic inadequacy in that language of having no category of noun phrases; it will persist even after that has been remedied. The problem we face here is that, unlike the case of *Sadie sleeps or Liz is-boring and Hank snores,* there are no syntactic or prosodic correlates of the evident semantic ambiguity in the sentence. Montague, like the Generative Semanticists, chose to set up multiple syntactic representations anyway in such cases and to retain the policy of assigning just one semantic value to each syntactic representation. As Cooper (1975) demonstrated, however, there is an alternative compatible with the Principle of Compositionality. One can have a single syntactic representation of *Some man loves every woman* and assign as the semantic value of each phrase a set of things very like the semantic values Montague assigns. Then the sentence turns out semantically ambigous in the appropriate way, while being unambiguous syntactically.