8.4. Empty Names and Empty Domains

A name is said to be empty if it denotes nothing, and we assumed at the beginning of Chapter 3 that names could not be empty. That is, we did not

[2] For some qualifications that seem to be needed here, see the appendix to this chapter.

allow name-letters to be interpreted as empty. Equally, we did not allow the domain of an interpretation to be empty. These decisions are connected, for if domains are allowed to be empty, then one must also allow names to be empty, as I now show.

Assume, for *reductio*, that a domain may be empty but that names must not be. Now in an empty domain it is clear that any formula beginning with an existential quantifier must be interpreted as false. Retaining the usual relation between $\exists$ and $\forall$, it then follows that any formula beginning with a universal quantifier must be interpreted as true. It is, as one says, 'vacuously' true. It should be observed that this ruling does fit with the semantics originally given for the quantifiers on p. 85. For if I is an interpretation with an empty domain, then there is no variant interpretation I_α for any name α. This is because a variant interpretation I_α must retain the same domain as I, but must also interpret α, and it cannot do both if the domain of I is empty but α cannot be interpreted as empty. But the semantics for $\exists$ says that $\exists\xi\varphi$ is to be interpreted as true in I iff there is a variant interpretation I_α which . . . , and we have just said that in an empty domain there is no such variant. Similarly, the semantics for $\forall$ says that $\forall\xi\varphi$ is to be interpreted as true in I iff for all variant interpretations I_α..., and this is vacuously the case if I has an empty domain. In an empty domain, then, $\forall\xi\varphi$ is always true and $\exists\xi\varphi$ is always false.

With this understanding it is clear that the sequent

$$\forall xFx \models Fa$$

remains correct, even if we include interpretations with an empty domain. For still there is no interpretation which makes $\forall xFx$ true and Fa false. To be sure, an interpretation with an empty domain will always make $\forall xFx$ true, but it cannot make Fa false, for since a has no interpretation on an empty domain, neither does Fa. In an entirely similar way the sequent

$$Fa \models \exists xFx$$

remains correct, for again there is no interpretation which makes Fa true and $\exists xFx$ false. By the Cut principle, you expect it to follow that the sequent

$$\forall xFx \models \exists xFx$$

must also be correct, and yet clearly this cannot be so, since the empty domain provides a counter-example. So one has to conclude that, in the situation envisaged, the Cut principle must fail. And in fact if you look back to the proof of that principle given on pp. 31–2 and p. 97, you will see that it requires an assumption which cannot be satisfied if we have empty domains but no empty names.

Not only does the Cut principle fail in this situation, but so also does what one might well call 'the oldest rule in the book', namely Modus Ponens. We need only a small modification of the example to make this point. By the same reasoning as before, the two sequents

$$\models Fa \vee \neg Fa$$
$$\models (Fa \vee \neg Fa) \rightarrow \exists x(Fx \vee \neg Fx)$$

remain correct, even if empty domains are permitted, as you may easily check. But their consequence by Modus Ponens, namely

$$\models \exists x(Fx \vee \neg Fx),$$

is no longer correct in this situation. These results seem to me to be wholly intolerable,[3] and I infer that any proposal that leads to them must therefore be rejected. In particular, then, we must reject the proposal that domains can be empty whereas names cannot.

The reason why this proposal led to an intolerable result was because it implies that there are formulae—i.e. those containing name-letters—which can only be interpreted on some of the permitted domains, but not on all. This situation is avoided if both names and domains may be interpreted as empty. For in that case a name *can* be interpreted on the empty domain, namely by interpreting it as denoting nothing. On this proposal the sequent

$$\models (Fa \vee \neg Fa) \rightarrow \exists x(Fx \vee \neg Fx)$$

is no longer correct. To see this clearly, take 'Fa' as abbreviating 'a exists'. Then on any domain $Fa \vee \neg Fa$ will be true, for if a is interpreted as denoting something in that domain, then Fa will be true, and if a is interpreted as not denoting—which it must be, if the domain is empty—then $\neg Fa$ will be true. But yet there is an interpretation in which $\exists x(Fx \vee \neg Fx)$ is false, namely one with an empty domain. In a similar way, both of the sequents

$$\forall xFx \models Fa$$
$$Fa \models \exists xFx$$

must be rejected as incorrect, if names may be empty as well as domains, as you are invited to work out for yourself. The two 'intolerable' results just mentioned are thus prevented on the new proposal. More generally you will find that the proof of the Cut principle originally given on p. 97 is now rescued.

[3] One has to admit that the results are *accepted* in Hodges (1977).

I conclude that if empty domains are permitted, then empty names must be permitted too. There is no equally strong argument for the converse conditional, that if empty names are permitted, then empty domains must be permitted too, but only a challenge: what motive might there be for allowing the one but not the other? This, however, brings us to what is obviously the main question: what positive motives are there for allowing either? Well, in each case the chief motive is that this represents a genuine possibility that should not be ignored. I offer three lines of argument for this conclusion. The first is a tricky line of argument, running into several problems and open to some quite plausible objections, which I shall not explore in any detail. So here I merely sketch the argument, since I think that it does have *some* force, but I do not pretend that what I say here adds up to a conclusive case. The second and third lines of argument are altogether more straightforward, and I think it is clear that they make a very strong case.

1. The first argument begins from the premisses (*a*) that logic is by tradition supposed to be an a priori science, i.e. one that needs no assistance from empirical enquiry, but also (*b*) that logic is used in the practical assessment of real arguments. Now, when we are aiming to test a real argument, the first thing that we need to do is to determine the domain of discourse of the argument, i.e. the domain over which the quantifiers employed are intended to range, and in practice one chooses different domains for different arguments. For example, one can often say that, for the purposes of this argument, only people need be included in the domain, or only items of furniture, or only cities, or whatever it may be. One can tell what the domain is supposed to be just by understanding what is being said at each point in the argument. But then, if the evaluation of the argument is to proceed a priori, we should not be relying upon our empirical knowledge when selecting a domain. Yet if domains have to be non-empty, then this cannot be avoided, for we cannot know a priori that there are people, or bits of furniture, or cities, and so on. A similar argument evidently applies to names. One cannot tell a priori whether a name that is being used as the name of a person really does denote a person, since one cannot tell a priori whether that alleged person does exist. Hence if our logic requires that an expression can be counted as a name only when it does denote something, then the logical testing of an argument cannot after all be an a priori process, as it was supposed to be.

This line of argument is open to various objections. One might try to avoid the point about domains by saying that logic never *requires* us to confine attention to this or that special domain, and we can always take the domain to be 'everything whatever' if we wish to. Then it can be added that, for logical purity, the domain should always be taken in this way, just because we *can* know a priori that there is at least *something*. But this reply raises many problems. For example, it could be argued (i) that logic does require one to confine attention to a special domain if the predicates (and functions) being considered are only defined on that special domain; (ii) that we cannot in fact make sense of the alleged domain of 'everything whatever'; (iii) that even if we can make sense of this domain, still our knowledge that it is non-empty cannot be a priori. I shall not pursue these problems any further.

Turning to the point about names, one might say that philosophers are now familiar with several theories of names, and some of them do seem to have the consequence that one cannot understand a name unless it does denote something. From this it may be inferred that one cannot understand a name without *knowing* that it denotes something, and hence that this knowledge must count as a priori. Again there are many problems. For example, (i) the inference from 'understanding requires that the name denotes' to 'understanding requires *knowing* that the name denotes' is surely questionable; (ii) whether it follows that this knowledge is a priori is also questionable, and must depend upon a detailed account of what a priori knowledge is; anyway (iii) the position invites this general response: if one cannot understand a genuine name without understanding that it is non-empty, then there can be only very few expressions that qualify as genuine names—perhaps only demonstratives such as 'this'. Once more, I shall not pursue these problems any further.

I remark finally about this line of argument that in any case one might wish to question the premiss upon which it is based. It certainly is part of the tradition that the study of logic contrasts with, say, the study of physics, on the ground that logic is a priori and physics is not. But there are plenty of philosophers nowadays who would call this tradition into question. I pass on, then, to my two other lines of argument, which do not invoke the notion of a priori knowledge.

2. The second argument retains the premiss that we are supposed to be able to apply our logic to test ordinary arguments for validity. Now, as I said in Section 1.2, an argument is a valid argument iff it is not possible for all its premisses to be true and its conclusion false. But in logic we do not study particular propositions, and particular arguments built from them, but proceed schematically, by working with formulae and sequents. A sequent is not itself a particular argument, but is rather a general pattern of argument, which actual arguments may exemplify. The idea is that if in our logic we can

show that a certain pattern of argument is a correct pattern, then it should follow that all arguments exemplifying that pattern are correct, i.e. valid, arguments. But how is this supposed to follow? Well, we count a pattern of argument as a correct pattern if there is no interpretation which makes all the premisses true and the conclusion false, and this is supposed to imply that in any argument of that pattern there is no possible situation in which all its premisses are true and its conclusion is false. The implication evidently depends upon the point that the interpretations considered in logic do exhaust all the possible situations for an actual argument. But this will not be so if interpretations are not permitted to have empty domains or empty names. For it *is* possible that there should have been nothing in the domain specified, whatever the specification, and it *is* possible that a named object should not have existed. These are genuine possibilities, even if we know (perhaps, in some cases, a priori) that they do not obtain in fact.

Someone may say: but why should we bother to take into account these possibilities which we know are not realized? I take this to be a proposal to amend the definition of validity for arguments, by building in a clause saying that a possibility can be ignored if it is known (or anyway, if it is *well* known?) that it is not realized. But this is quite contrary to the spirit of logic. We can see this clearly if we look back to an interesting episode in the development of the subject. It is well known that Aristotle's system of syllogistic logic accepted as correct some laws which nowadays we reject, for example

$\models$ (Some *F*s are *G*) or (Some *F*s are not *G*).

The explanation is that Aristotle was failing to take into account the possibility of there being no *F*s at all. So we say today that several of the argument-patterns which he accepted as correct are not actually correct, for they need the extra premiss

There are some *F*s.

We may concede to him that very often this extra premiss would state only what was well known both to the arguer and to his audience, and so it would be unsurprising if in practice it was often omitted. After all, in practice we often do omit all kinds of premiss as too obvious to need an explicit statement. But still we should insist that logic alone cannot certify the argument to be correct until the extra premiss is explicitly put in. So too, it seems to me, with the possibility of a whole domain of quantification being empty, and of a named object failing to exist. What is still today the standard logic ignores these possibilities, but that means that it is sometimes *mistaken* in which arguments it counts as valid. For some of these arguments require extra premisses, stating that a named object does exist, or that at least something exists, if they are to be valid in the proper sense. Again we may concede that in practice these extra premisses very often go without saying, since they are well known to all participants. But that should not prevent us from insisting that pure logic can certify such an argument to be valid only if the missing premisses are explicitly put in.

3. My third argument no longer concerns the practical application of logic in testing ordinary arguments, but the role that elementary logic has as forming the basis on which to build more advanced logics. Here I briefly consider just one such more advanced logic, namely modal logic, which studies the two sentence-functors 'it is possible that' and 'it is necessary that'. These are abbreviated to '$\Diamond$' and '$\Box$' respectively. One sets out such a logic by first assuming the appropriate elementary logic, and then adding more rules or axioms to deal with the modal sentence-functors. A standard rule for this purpose, adopted in many systems, is the so-called rule of necessitation, stating that what can be proved is necessary:

If $\vdash \varphi$ then $\vdash \Box\varphi$.

But this rule is clearly incorrect if as our underlying elementary logic we take the logic studied up to this point, which does not permit empty names. For in this logic we have, for any name a,

$$\vdash \exists x(x=a).$$

But we do *not* want

$$\vdash \Box\exists x(x=a).$$

For even though Margaret Thatcher does exist—and even if the name 'Margaret Thatcher', used as we use it, would not have existed unless Margaret Thatcher had existed—still it is not a necessary truth that Margaret Thatcher exists. On the contrary, it is evidently possible that she should not have done. (For she would not have existed if her mother had died at the age of 10, and that is something that might have happened). An entirely similar point holds about empty domains. In the logic studied up to now we have, e.g.

$$\vdash \exists x(x=x).$$

But we do *not* want

$$\vdash \Box\exists x(x=x).$$

For it is a possibility that nothing should have existed at all.

It may be replied to this argument that the rule of necessitation is not sacrosanct, and there is in principle no reason why we should not retain our existing elementary logic and modify or abandon this rule. That is, no doubt, an avenue that one might explore. But it is surely a more straightforward course to modify the elementary logic so that only necessary truths can be proved in it. The rest of this chapter presents such a modification.

EXERCISES

8.4.1. Look back to the proof in section 3.7 that for every formula there is an equivalent formula in PNF.
(*a*) Explain where that proof breaks down if empty domains are permitted.
(*b*) Show that this breakdown cannot be repaired, i.e. that there is a formula which has no equivalent in PNF, if empty domains are permitted.
(*c*) Do you think this provides a good argument for saying that empty domains should not be permitted?

8.4.2. If you know of any philosophical arguments in favour of the dictum 'existence is not a predicate', consider whether those arguments show that the traditional logic is right to discount the possibility of a named object not existing.

8.5. Extensionality Reconsidered

In Section 3.1 I introduced two principles about names which underlie the whole treatment of names in the traditional logic. One was that a name always has a denotation, and I have just been arguing that this principle should be rejected. The other was the principle of extensionality, that if two different names denote the same object, then they behave as if they were the same name, i.e. either may be substituted for the other in any context. This is reflected, of course, in the adoption of Leibniz's law,

$$a = b \models (Fa \leftrightarrow Fb),$$

as a principle governing identity. I noted at the time that in a natural language there would be many occurrences of names that seemed to conflict with this principle, but said that for the purposes of elementary logic these must just be set aside. For we cannot recognize anything as an occurrence of a name unless the principle of extensionality applies to it. But if we decide to jettison the first principle, that names must always have denotations, then what becomes of the second?

Well, since there do seem to be many exceptions to this second principle, one might wish to jettison it as well. And there would be good reason to do so if we were trying to develop a more advanced logic concerning what is necessary or possible, or what is knowable a priori, or simply what is known or believed or held probable or something of the sort. For in such contexts as these the principle of extensionality very frequently seems to fail. But that is not our present task. At the moment we are simply considering the ordinary and straightforward areas of language which one usually regards as subject only to elementary logic, and not needing an advanced treatment. Is there any need to say that in these simple contexts extensionality must be abandoned once empty names are permitted?

There is not. On the contrary, there is a great need to retain the principle so far as possible, for, as I pointed out (p. 74), it is built into the ordinary semantics for elementary logic. For example, when a name-letter is interpreted as denoting something, then all that we provide by way of an interpretation is the object denoted. Similarly, all that we provide by way of interpretation for a one-place predicate-letter is the set of objects that it is true of. There is nothing in this simple apparatus that could explain how a predicate might be 'true of' an object under one name but not under another, and it would clearly be going beyond the confines of elementary logic if we tried to introduce a more complex apparatus. I conclude that extensionality must be retained when we are dealing with names that do denote, so the problem is: how are we to explain extensionality for names that do not denote? I think the answer is quite straightforward. If the truth-value of a predication depends only on *what* the name denotes, so that it must remain the same for any other name denoting the same thing, then when it turns out that the name denotes nothing, that fact *itself* must determine the truth-value. So the truth-value will remain the same for any other name that also denotes nothing.

To illustrate, consider the simple predicate 'x is a horse'. It is undeniable that if in place of 'x' we have a name that denotes something, then whether the whole is true or not depends only on whether the thing denoted is a horse. And it cannot possibly happen that two names denote the same thing, but one of them denotes a horse and the other does not. So the principle of extensionality is certainly satisfied in this case. It is also satisfied with names that denote nothing, if we say—as seems to me very reasonable—that only what exists can be a horse. Thus it is not true that Jupiter is a horse, principally because Jupiter does not exist, though in this case one may wish to add that even if he did exist, he still would not be a horse. (Here is an interesting question for philosophers: how do you know that?) Equally, it is not true that Pegasus is a horse, and again this is because Pegasus does not exist (and never did). Here one is tempted to say that if Pegasus had existed he would have been a horse. By the same token, he would have been winged. So there would have been a winged horse. But as things are, there are no winged horses, since Pegasus is not a horse. And this is simply because Pegasus does not exist (now, or at any other time).

Two comments may be made at once. First, there are people who will protest that 'Pegasus is a horse' should be accepted as true, and one must admit that we do often talk in this way. When our discourse is about fictional characters, it appears that we take a domain of quantification that includes these characters, and we count it as true that P if the relevant story says (or implies) that P, and as false that P if the relevant story says (or implies) that $\neg P$. From a logical point of view, however, one cannot take this proposal seriously, since it must lead to a breakdown in elementary logical laws. Thus $P \vee \neg P$ will not be true when the relevant story says nothing either way, and, worse, $P \wedge \neg P$ *will* be true when the relevant story is inconsistent. So we would do better to say that this way of talking is really a shorthand. We talk as if we took it to be true that P when all that we really mean—or should mean—is that it is true that *it is said in the story that P*. To apply this to the example in hand, we should continue to insist that

Pegasus is a horse

is *not* true, but we add that there is a related statement which *is* true, namely

It is said in Greek mythology that Pegasus is a horse.

The two statements are, of course, different statements, and the first would follow from the second only if whatever is said in Greek mythology is true. But the fact is that most of what is said in Greek mythology is not true, and this includes the claim 'Pegasus is a horse'.

Second, there are people who, when persuaded that 'Pegasus is a horse' is not true, think that the same should therefore apply to 'Pegasus is not a horse'. Their thought is that *neither* of these can be true if Pegasus does not exist. But I see no reason to agree. For negation is so defined that '$\neg P$' counts as true in any situation in which 'P' is not true, and so in the present situation '$\neg$(Pegasus is a horse)' *is* true. If this is rejected, then certainly the familiar logic can no longer be upheld, and we would apparently need a 'third truth-value', i.e. 'neither true nor false'. But such a reaction is surely too extreme. Observe, first, that we must certainly reject the perfectly general principle that *nothing* can be true of *a* unless *a* exists, for when *a* does not exist it is quite clear that '¬(*a* exists)' *is* true of *a*. It follows that when '*Fa*' implies '*a* exists', and the truth is that ¬(*a* exists), then also the truth is that ¬*Fa*. People who are still uneasy about this may perhaps be placated by this suggestion: perhaps they are understanding 'Pegasus is not a horse' not as the negation of 'Pegasus is a horse', but as equivalent to

(Pegasus exists) ∧ ¬(Pegasus is a horse).

Certainly, *this* is no more true than 'Pegasus is a horse'.[4]

To return to the original question concerning extensionality, I propose that this principle be preserved in an elementary logic which admits empty names by the ruling that all empty names are to behave alike, i.e. that substituting any one for any other will always leave truth-values unchanged. The suggested elucidation of this is that where '*F*' represent an *atomic* predicate, such as '... is a horse', then '*Fa*' will be false whenever '*a*' denotes nothing. The truth-values of more complex sentences containing '*a*' will then be determined in the usual way by the truth-values of their atomic components; in particular, if '*Fa*' is false when '*a*' denotes nothing, then '¬*Fa*' will be true when '*a*' denotes nothing. I observe here that this gives the right result when '*Fa*' is '*a* exists'. This is an atomic statement, so by the suggested ruling, when '*a*' denotes nothing we shall have '*a* exists' as false and '¬(*a* exists)' as true. This is evidently as it should be.

This ruling has an effect upon what sentences of ordinary language we can accept, for logical purposes, as made up from a predicate and a name. For example, consider a sentence

John is painting a picture of *a*.

It is quite easy to see this sentence as satisfying the principle of extensionality when the name '*a*' does denote something, but we still have a problem when it does not. Previously we had to say that we cannot count this as a sentence of the form '*Fa*' when '*a*' fails to refer, since our stipulation was that all names must refer. Now we do not rule it out on this ground, but we must still rule it out nevertheless. For 'Pegasus' refers to nothing, and so does 'Jupiter', so if we accept

4 Philosophers will note at this point that I am making no distinction between on the one hand saying or implying that Pegasus exists, and on the other hand *presupposing* this point. That is a fair comment, but one which I do not propose to discuss. Presupposition can have no place in elementary logic.

John is painting a picture of Pegasus

as containing the name 'Pegasus', then we must also accept

John is painting a picture of Jupiter

as obtained from it by substituting one empty name for another. But then our extensionality principle would require us to say that the two sentences must have the same truth-values, which is wholly paradoxical, since it is clear that a picture of Pegasus is not at all the same thing as a picture of Jupiter. The result then is that for the purposes of elementary logic we still cannot accept 'John is painting a picture of Pegasus' as made up from a name 'Pegasus' and a predicate 'John is painting a picture of . . .'.

This situation is common. Where previously we had to say that an apparent example of name-plus-predicate structure could not be accepted at face value, just because the name was empty, so now we quite often have to reach the same conclusion, but on the different ground that empty names do not all behave alike in that context. But this does not happen always, and there are some sentences which we can now recognize as having the form '*Fa*' but could not have done before. The simplest and most prominent examples are the sentences

$\neg$(a exists).

Previously we had to say that no such sentence was of the form '*Fa*', because we required that, for any admissible sentence '*Fa*',

$$Fa \vDash \exists xFx.$$

Naturally, this does not hold for '$\neg$(a exists)' in place of '*Fa*'. But with our revised conception it no longer has to.

Let us move on, then, to consider just what rules of inference do hold on the revised conception.

EXERCISE

8.5.1. Consider the sentence

John is writing a story about King Arthur.

Do we have to say that this sentence means one thing if King Arthur did exist, but a different thing if there was no such person?

8.6. **Towards a Universally Free Logic**

A 'free' logic is one in which names are permitted to be empty. A 'universally free' logic is one in which the domain of an interpretation may also be empty, and our object now is to formulate such a system. In this section I approach the topic by asking how the familiar rules of inference need to be modified for this purpose. For definiteness, I shall concentrate attention upon the rules required for a tableau system of proof.

As I argued on pp. 357–8, the rules for the truth-functors are not affected by the new view of names and domains, so we may turn at once to the quantifier rules. The familiar tableau rules are

$$\begin{array}{cccc} \forall\xi\varphi & \neg\forall\xi\varphi & \exists\xi\varphi & \neg\exists\xi\varphi \\ | & | & | & | \\ \varphi(\tau/\xi) & \neg\varphi(\alpha/\xi) & \varphi(\alpha/\xi) & \neg\varphi(\tau/\xi) \end{array}$$

provided α is new

But the intended interpretation is now that the quantifiers range only over existing things (as before), whereas the terms are not so restricted, and this means that each rule requires a modification. First, the two outer rules need to be weakened. For, if τ is a term that fails to denote, then from the premiss that all *existing* things satisfy φ it will not follow that τ satisfies φ. (For a clear counter-example take $\varphi(\tau/\xi)$ as 'τ exists'.) So here we need to add the extra premiss that τ exists, if the inference is to remain sound. The same evidently applies to the rule for $\neg\exists$. By contrast, the two inner rules can be strengthened. For example, the premiss to the $\exists$ rule tells us that there *exists* something satisfying φ, and we argue as before that we can therefore introduce a name α for that thing, provided that the name is a new name, i.e. one that has not already been used for anything else. It then follows as before that α satisfies φ, but now we can *also* add something further, namely that α exists. Abbreviating 'α exists' to '$E!\alpha$', our four rules therefore need modifying in this way:

$$\begin{array}{cccc} \forall\xi\varphi & \neg\forall\xi\varphi & \exists\xi\varphi & \neg\exists\xi\varphi \\ E!\tau & | & | & E!\tau \\ | & E!\alpha & E!\alpha & | \\ \varphi(\tau/\xi) & \neg\varphi(\alpha/\xi) & \varphi(\alpha/\xi) & \neg\varphi(\tau/\xi) \end{array}$$

provided α is new

So far this is all very straightforward, and all free logics are in agreement. It is easy to see that the quantifier rules do need modifying in these ways, and that the modified rules are sound under the revised conception. Hence we cannot prove from them any sequents which must now be counted as incorrect, such as

$$\forall xFx \models \exists xFx$$

or

$$\models \exists x(Fx \lor \neg Fx)$$

or indeed

$$\models \exists xE!x.$$

But we can prove some new sequents which are correct on the new conception, such as

$$\models \forall xE!x.$$

The next thing to ask, therefore, is whether we can prove *all* such sequents, i.e. whether the new quantifier rules are complete for the new conception.

The answer to this depends upon what exactly the new conception is. In particular, I have argued in the last section that the principle of extensionality should be extended to empty names by requiring that all empty names behave alike, and this is a thesis that can certainly be formulated in the vocabulary now being used, thus

$$\text{EXT: } \neg E!\tau_1, \neg E!\tau_2 \models \varphi(\tau_1/\xi) \leftrightarrow \varphi(\tau_2/\xi)$$

('EXT' is for 'extensionality'.) But we certainly cannot prove this principle from the quantifier rules already stated. Now some free logics do not adopt this principle EXT, and they may count the four quantifier rules already given as complete. But I have argued that EXT should be adopted, and since we cannot prove it from what we have already, on my conception these four rules are not complete. Suppose, then, that EXT is added as a new rule of inference. Will that give us a complete system? Well, in a sense, yes.[5] But our present rules use *E!* as a primitive and undefined predicate, because they do not yet say anything about identity. Once we add identity, however, *E!* will surely be definable, for we expect it to be true that

$$E!\tau \dashv\models \exists\xi(\xi = \tau).$$

So before I come back to the question of completeness, I now proceed to consider what rules we ought to have for identity in a free logic. (On this question there is no general agreement amongst logicians.)

[5] But also, in a sense, no. The position will be clarified in the following section.

The principle of extensionality for names that do denote objects holds under the new conception just as much as it did under the old. So we may lay it down that

$$\text{LL: } \tau_1 = \tau_2 \vDash \varphi(\tau_1/\xi) \leftrightarrow \varphi(\tau_2/\xi)$$

('LL' is for 'Leibniz's law'). On this there is universal agreement, so far as I know. But it is not so clear whether we should retain the other principle for identity,

$$\vDash \tau = \tau$$

for *all* terms τ whatever, or whether we should say that only existent things can be identical with anything, and hence that only they can be self-identical.

We should observe first that EXT has this implication: if even non-existent things are still counted as self-identical, then all non-existent things must be counted as identical with one another. For, as an instance of EXT, we have

$$\neg E!a, \neg E!b \vDash (a{=}a \leftrightarrow a{=}b).$$

If, then, $a{=}a$ holds for all names a whatever, whether or not a exists, we can infer

$$\neg E!a, \neg E!b \vDash a{=}b.$$

This apparently says that there is at most one non-existent thing. As such, it is a principle well suited to the approach whereby a name that appears to denote nothing is always treated, despite appearances, as denoting something —either some arbitrarily chosen and familiar object, such as the number 0, or perhaps a specially invented object called 'the null object'. (This 'null object' must then be a member of *all* domains, even those that we think of as empty, since it must be possible to interpret a name-letter on any domain.) But such an approach is hardly attractive. I have already noted (pp. 338–9) that we obtain unwanted truths if we suppose that what appears to name nothing does actually name an arbitrarily chosen but familiar object. We do not get this consequence with 'the null object', since it is not a familiar object already figuring in familiar truths. On the contrary, it is a wholly unfamiliar object, invented just to be the thing that all otherwise empty names denote. As such, one must admit that it is technically convenient, and a neat way of upholding the principle that all empty names behave alike. But, at the same time, it is sheer fantasy, and we do not actually *need* any such fantasy. We can perfectly well say that a name may be interpreted as denoting nothing, without having to suppose that this 'nothing' is really something, but a strange thing.

If we do dispense with the fantasy, then the suggestion that all non-existent things are identical with one another will get no support in this way, and I think that it gets no support in any other way either. But there is an argument against it. For I have said that a good general principle is that when F represents an atomic predicate, and a a name that denotes nothing, then Fa is to be false. This principle seems to me to work very well for a wide range of simple atomic predicates. But then we have only to add that identity is to be regarded as a simple atomic predicate, and we have a definite ruling on the present issue: if a does not exist, then $a=b$ is always to be false, whatever b may be; hence as a special case $a=a$ is false too.[6]

Now the sequent

$$b=a \models a=a$$

is provable just from LL by itself. Hence by the quantifier rules we already have a proof of

$$\exists x(x=a) \models a=a.$$

The decision that $a=a$ is to be true only when a exists now allows us to affirm the converse

$$(*)\ a=a \models \exists x(x=a).$$

It follows that either $a=a$ or $\exists x(x=a)$ would do equally well as our analysis of $E!a$. It turns out, however, that there is an advantage in choosing the first. For if we write $\tau=\tau$ in place of $E!\tau$ (and $\alpha=\alpha$ in place of $E!\alpha$) in the quantifier rules as formulated on p. 360, then it turns out that (*) is deducible, whereas if we write $\exists\xi(\xi=\tau)$ instead, then (*) will be needed as a separate postulate. I therefore adopt this decision. We shall have as a definition

$$E!\tau \quad \text{for} \quad \tau=\tau.$$

Our basic rules will then be the four quantifier rules of p. 360, but with $E!$ eliminated in favour of its definition, and in addition just EXT and LL as further rules. In the context of a tableau system of proof, these are formulated as

[6] One *might* claim that identity should *not* be regarded as a simple atomic predicate. In a *second*-order logic it may be defined thus:

$$a=b \quad \text{for} \quad \forall F(Fa \leftrightarrow Fb).$$

On this definition it is not atomic, and $a=a$ is always true, whether or not a exists. (Similarly, $a=b$ is true if neither a nor b exists.)

$$\begin{array}{c} \tau_1 \neq \tau_1 \\ \tau_2 \neq \tau_2 \\ \varphi(\tau_1/\xi) \\ | \\ \varphi(\tau_2/\xi) \end{array} \qquad \begin{array}{c} \tau_1 \neq \tau_1 \\ \tau_2 \neq \tau_2 \\ \varphi(\tau_2/\xi) \\ | \\ \varphi(\tau_1/\xi) \end{array} \qquad \begin{array}{c} \tau_1 = \tau_2 \\ \varphi(\tau_1/\xi) \\ | \\ \varphi(\tau_2/\xi) \end{array} \qquad \begin{array}{c} \tau_1 = \tau_2 \\ \varphi(\tau_2/\xi) \\ | \\ \varphi(\tau_1/\xi) \end{array}$$

This completes our combined theory for quantifiers and identity together.

I remark here that some free logics adopt the principle that, if Φ^n is any *atomic* predicate-letter, then for $1 \leqslant i \leqslant n$

$$\Phi^n(\tau_1,...,\tau_n) \models E!\tau_i.$$

I think, however, that this is due to a confusion. I have already said that I think this is a very reasonable thesis where Φ^n is an atomic *predicate*, but predicates are not the same as *predicate-letters*. On the contrary, the role of the predicate-letter is to take the place of *all* kinds of predicates, and not only atomic ones. That is why the principle of uniform substitution for predicate-letters is a correct principle.

A similar principle that is sometimes proposed applies this idea to function-letters. The suggestion is that if θ^n is any n-place function-letter, then we should have, for $1 \leqslant i \leqslant n$,

$$E!\theta^n(\tau_1,...,\tau_n) \models E!\tau_i.$$

But I think that this suggestion too should be rejected. No doubt many familiar functions do obey the proposed condition that they are defined—i.e. have a value which exists—only for arguments which themselves exist. But there seems to be no reason to insist that all functions whatever must be like this. For example, it is often useful to introduce, for any predicate, its corresponding 'characteristic function'. The definition takes this form: given a predicate F—a one-place predicate, for simplicity—we introduce the corresponding function f defined so that

$$f(\tau) = 1 \ \text{ if } \ F(\tau)$$
$$f(\tau) = 0 \ \text{ if } \ \neg F(\tau).$$

Such a function must provide a counter-example to the proposed principle. For, as we have said, even where τ does not exist, still either $F(\tau)$ or $\neg F(\tau)$ will be true, so in any case $f(\tau)$ will be defined. I shall not, then, add any further principles of this sort.

In fact I prefer not to add function-letters at all, since there is nothing that can be laid down as a general principle to say when a function is or is not defined for a given argument. So the system that I have been considering so far is one in which the only closed terms are simple name-letters, and there is therefore no practical distinction between the letters τ and α. But I shall now proceed to add definite descriptions as further terms. As I have noted, functions are a special case of descriptions, so if descriptions are available, then we do not need to make any further provision for functions. And because a definite description necessarily has an internal structure, it is easy to say when there is such a thing as it describes, namely when the description is uniquely satisfied.

Since our system allows names to be empty, the chief objection to treating definite descriptions as (complex) names has now disappeared, and I shall therefore treat them in this way. Accordingly, given any formula φ containing a free occurrence of the variable ξ, the expression $(\iota\xi{:}\varphi)$ will be a term. Since all occurrences of ξ in this term are bound (by the initial prefix $\iota\xi$, if they are not already bound in φ), the term will be a closed term provided that there are no other variables free in φ. In that case, the rules that have been stated as applying to all terms τ will apply to it. If we wish to allow our rules to apply also to open formulae, which may contain open terms, then as before (p. 337) we need to make sure that the substitution-notation is so explained that all occurrences of variables that are free in τ remain free in $\varphi(\tau/\xi)$. Once this is done, the extension presents no further problem. If we wish to allow for vacuous occurrences of the prefix $\iota\xi$, i.e. occurrences in which it is attached to a formula φ containing no free occurrences of ξ, then this is harmless. In nearly all cases it will lead to a term $(\iota\xi{:}\varphi)$ that does not denote, though there is one exception (Exercise 8.6.4). Finally, the obvious principle for when a definite description has a denotation, is this:

$$E!(\iota\xi{:}\varphi) \;\leftrightarrow\; \exists\zeta\forall\xi(\varphi \leftrightarrow \zeta{=}\xi).$$

A stronger principle from which this follows, and which I adopt as an axiom-schema governing descriptions, is this:

$$\text{(PD)}\quad \forall\zeta(\zeta{=}(\iota\xi{:}\varphi) \leftrightarrow \forall\xi(\varphi \leftrightarrow \zeta{=}\xi)).$$

('PD' is for 'principle of descriptions'.) Considered as a rule of inference for the tableau system, the rule is that any instance of this principle may be added to the root of any tableau.

8.8. Appendix: A Note on Names, Descriptions, and Scopes

A standard assumption in logic is that names do not have effective scopes. For example, it does not make any difference whether we see the formula $\neg Fa$ as obtained by first supplying the name a as subject to the predicate F, and then negating the result, or as obtained by first negating the predicate F and then supplying the name a as subject to this complex predicate. We may look at it either way. On Russell's analysis of definite descriptions, this gives us a clear contrast between names and descriptions, for it is a feature of his analysis that descriptions must have scopes. To use the same example, if we see the formula $\neg F(\iota x{:}Gx)$ as got by first supplying $(\iota x{:}Gx)$ as subject to the predicate F, and then negating the result, the whole will be true if $(\iota x{:}Gx)$ fails to exist and F represents an atomic predicate. But if we see it as got by first negating the predicate F, and then supplying $(\iota x{:}Gx)$ as subject to this complex predicate, we get the opposite result. This is the difference between

$$(\text{I}x{:}Fx)(\neg Gx) \quad \text{and} \quad \neg(\text{I}x{:}Fx)(Gx).$$

This point suggests the following thought. If it is really true that definite descriptions have scopes whereas names do not, then Russell must be right to claim that definite descriptions are not names. If, however, this is not really true, then it does no harm to treat descriptions as complex names, which is what the system B does.

Now, provided that names are allowed to be empty, as descriptions evidently can be, this question cannot be decided at the level of elementary logic. For, as I have pointed out (Exercise 8.3.2), in elementary logic when we assign the scope of a description in one way rather than another, this can make a difference only if the description is empty, and the difference that it then makes is just that some versions imply that the description is not empty whereas others do not. (Thus, concerning the examples of the last paragraph, the first implies that $(\iota x{:}Gx)$ is not empty, and the second does not; on the contrary it is true if F is atomic and $(\iota x{:}Gx)$ is empty.) But we can get exactly the same effect with names, not by assigning them scopes, but by including or excluding explicitly existential clauses. For example, the distinction just noted for a description $(\iota x{:}Gx)$ can be reproduced for a name a as the distinction between

$$E!a \wedge \neg Fa \quad \text{and} \quad \neg(E!a \wedge Fa).$$

Thus all the work that could be done, in elementary logic, by making scope-distinctions either for names or for descriptions, can equally well be done instead by adding explicitly existential clauses at appropriate points in the formula. I conclude that, at the level of elementary logic, there is no call to assign scopes either to names or to descriptions.

It may be replied that, when we move to a more complex level of discourse, where we cannot make do just with the resources of elementary logic, the advantage of

Russell's analysis of descriptions becomes clear. For example, Kripke (1980: 48–9) has argued that, where *a* is a name, the sentence

It might not have been that $a=a$

is unambiguous, and always false. But if in place of the name *a* we take a description $(\iota x{:}Fx)$, then—he says—we must make a distinction. The proposition

It might have been that $(\mathrm{I}x{:}Fx)\neg(x=x)$

is always false: there is no possibility of there being one and only thing which is *F* but not self-identical. But, in contrast, it is very often true to say

$(\mathrm{I}x{:}Fx)$(It might have been that $\neg(\mathrm{I}y{:}Fy)(x=y)$).

This is true because it means that, concerning the one and only thing that is *F*, it might not have been the one and only thing that is *F*—either because it might not have been *F* or because something else as well might have been *F*. But—he claims—there is no similar point to be made about a name *a*: we cannot say that, concerning the thing that is *a*, it might not have been *a*.

For the sake of argument, let us grant that Kripke is essentially right on this point. Nevertheless, one might still find the point unimpressive, since there are many other non-extensional contexts where—at least at first sight—scope-distinctions seem to be needed just as much for names as for descriptions. To adapt an example of Quine's (1960: 141–56), consider the sentence

Ralph believes that the man in the brown hat is a spy.

This may be taken in two ways. It can be understood as saying that Ralph believes to be true what is said by the whole sentence 'the man in the brown hat is a spy', or as saying that Ralph believes to be *true of* a certain person what is expressed by the predicate 'that he is a spy', where the person in question is in fact the one man here in a brown hat, though Ralph may be unaware of this fact. The distinction is that in the first case the words 'the man in the brown hat' are taken as part of the report of what Ralph believes, whereas in the second case they are the speaker's way of referring to a particular person, which need not also be Ralph's way of referring to him. In the jargon, the sentence is said to be understood *de dicto* in the first case and *de re* in the second. Now at first sight it is tempting to say that the distinction is one of scope. In the first case we have

Ralph believes that $(\mathrm{I}x{:}x$ is wearing a brown hat) (x is a spy)

and in the second case we have

$(\mathrm{I}x{:}x$ is wearing a brown hat) Ralph believes that (x is a spy).

But here one must notice that exactly the same ambiguity occurs when we have a name in place of the description, as in

Ralph believes that Bernard J. Ortcutt is a spy.

Again, the name 'Bernard J. Ortcutt' may be taken as reporting part of the content of Ralph's belief, or it may be taken as the speaker's way of telling us who is the object of Ralph's belief. Should we conclude, then, that in contexts such as these both names and descriptions should be assigned scopes?

Quine, for one, would not wish to look at it in this way. On his suggestion the ambiguity is better viewed, not as a question of the scope of the name or description, but as an ambiguity in the prefix 'Ralph believes that'. We can construe this prefix as an operator that forms a sentence from a sentence, or we can construe it as an operator that forms a sentence from a name and a predicate taken separately, where only the predicate represents what is believed, and the name is used to say what it is believed of. In the second case, then, the more perspicuous rendering is

Ralph believes, of Bernard J. Ortcutt, that he is a spy.

As Quine insists, we must be able to understand belief-sentences *both* in the one way *and* in the other. Belief is what he calls a 'multigrade relation', relating a believer either to a sentence, or to an object and a one-place predicate, or perhaps to two objects and a two-place predicate, and so on.

It is a question for philosophical disputation whether Quine's way of looking at these sentences is better or worse than the way which assigns scope to referring expressions, or whether the apparent disagreement between these two approaches is one that disappears on a closer analysis. I shall not here take this dispute any further. But in any case we can say that contexts of this kind provide no motive for *distinguishing* names from descriptions. Bearing this in mind, let us look back once more to Kripke's case for saying that descriptions do have scopes whereas names do not. Clearly Quine would wish to distinguish between

(1) It might have been that *a* was not *a*.
(2) It might have been, concerning *a*, that it was not *a*.

The first takes 'it might have been that' to be operating on a whole sentence, whereas the second takes it to be operating on a name and a predicate taken separately. Now for the sake of argument we may agree with Kripke that (1) is always false, whereas (2) may be true where '*a*' is a description, but not where '*a*' is a name.[8] But this, we may suggest, is a peculiar feature of the way that names interact with modal operators such as 'it might have been that', and is to be explained by the fact that names are what Kripke calls 'rigid designators'. This means (roughly) that names continue to designate the same thing as we shift our attention from one possible situation to another, whereas most descriptions do not. But this point by itself would not prevent us from saying that descriptions may be treated as complex names, for it does not in any way imply that descriptions have scopes whereas names do not. On the contrary, the only 'scope-distinction' that is here envisaged is a distinction in how the prefix 'it might have been that' is to be understood, and this has no tendency to show that definite descriptions are not complex names.

[8] Actually, one could perfectly well claim that (2) will be true wherever '*a*' is a name that might have been empty. To avoid this objection, change the example to

It might have been, concerning *a*, that it existed but was not *a*.