THE MODAL REALISM OF DAVID LEWIS

THE MODAL REALISM

OF

DAVID LEWIS

By

ANDREW SULLIVAN, B.A.

A Thesis

Submitted to the School of Graduate Studies

in Partial Fulfilment of the Requirements

for the Degree

Master of Arts

McMaster University

MASTER OF ARTS (1996) McMaster University
(Philosophy) Hamilton, Ontario

TITLE: The modal realism of David Lewis

AUTHOR: Andrew Sullivan, B.A. (University of Ottawa)

SUPERVISOR: Doctor M. Vorobej

NUMBER OF PAGES: vii, 128

ABSTRACT

It seems everyone who encounters David Lewis's modal realism finds it utterly strange. I attempt to show that it is strange because it expands our ontology unnecessarily.

I begin by situating Lewis in the tradition of understanding modality with the help of possible worlds. I go on to indicate the internal tensions which arise under Lewis's view when we consider some kinds of perfectly normal modal discourse. I try to show that the problem comes from an understanding of existence which is quite common in analytic philosophy. I then suggest a way of understanding modal language which finds its inspiration in the writings of St Thomas Aquinas, Ludwig Wittgenstein, and Donald Davidson.

ACKNOWLEDGMENTS

Had it not been for Dr Mark Vorobej, this thesis would never have seen completion. Fortunately, Dr Vorobej actually supervised my work. I am grateful for his tremendous help and his well of patience. I am also thankful for the insightful comments of Dr Nicholas Griffin. I am also forever grateful for the willing ear of Dr Jill LeBlanc.

TABLE OF CONTENTS

LIST OF ABBREVIATIONS

"Counterpart Theory"	Lewis, "Counterpart Theory and Quantified Modal Logic"
CSM III	Descartes, *The Philosophical Writings of Descartes*
"OD"	Russell, "On Denoting"
PoM	Russell, *Principles of Mathematics*
PT	Quine, *Pursuit of Truth*
PW	Lewis, *On the Plurality of Worlds*
"RM"	Quine, "Reference and Modality"
"SC"	Kripke, "Semantical Considerations on Modal Logic
ST	St Thomas, *Summa Theologica*

For full citations, please see the bibliography.

Chapter 1

Possibility and Worlds

1.1 Modality and possible worlds

"I could have written about something else." Nothing is more obvious than the truth of that sentence. We -- or, at any rate, most of us -- believe that the world could be different in countless ways. I could have studied engineering, or the law; I could have had black hair; donkeys could have been able to talk, and elephants, able to fly. Our statements to the effect that some state of affairs could be different are simply evidence that we believe in alternative possibilities.

What is strange about this ubiquitous notion of possibility is how difficult it is to say what it is, exactly. We can provide lots of examples; but that does not offer us a clarification of the concept, 'possibility'. It is not terribly enlightening to say, "A state of affairs, ϕ, is possible when a statement to that effect, 'ϕ', could be true." For that does not get us any closer to an understanding of what it would mean for a statement of the form 'possibly-ϕ' to be true. And if we are puzzled about the notion of possibility, we will want to know when such a sentence is true. We believe such statements all the time: it seems perfectly obvious that I might have studied engineering, for example. What we would like, then, is a way of understanding when these modal sentences are true.

Perhaps the way to understanding is via the notion of necessity. Something is necessarily true when it cannot be false. But all that says is that something is necessarily true when it is not possible for it to be false. Nevertheless, it seems obvious that there are ways the world could *not* be. It seems obviously impossible that an object be both round and square. It seems equally impossible that an object not be identical to itself, or that 3+5 be equal to anything other than 8. But we have not yet offered any criteria for telling which modal sentences are true, and which are false.

Perhaps we can best understand modal sentences as being about states of affairs. On such a view, the usual declarative, non-modal sentences are just about the state of affairs as it actually obtains. "The cat is on the mat," is true only if there is a cat actually on the mat. Similarly, modal sentences are about the actual state of affairs, *and also* about other, possible states of affairs. "I could have written about something else," is true only if, under some state of affairs, I *do* write about something else. We need not worry about describing the nature of these states of affairs just now; all we need to understand is that, just like there are actual states of affairs, there are possible states of affairs, too. For short, we can call these possible states of affairs (or, more precisely, some sets of these possible states of affairs) "possible worlds". So, a statement of the form 'possibly-ϕ' is true just in case there is some possible world where 'ϕ' is true. A necessary truth is one which cannot be false at any world. Now all we

have to do is explain what the worlds are like. That task, as it turns out, is easier said than done.

1.2 Leibniz's worlds

The notion of possibility as truth-in-a-world is often attributed to Leibniz. Leibniz wanted to understand what it means to say that God could have made things differently. It seems that God could have made the world differently, given His nature; for, Leibniz says,

> 37. And as all this differentiation involves only other prior or more differentiated contingent things, all of which need a similar analysis to explain them, we are no further advanced: and the sufficient or ultimate reason must be outside the succession or *series* of this differentiation of contingent things, however infinite it may be.
>
> 38. This is why the ultimate reason of things must lie in a necessary substance, in which the differentiation of the changes only exists eminently as in their source; and this is what we call *God*.
>
> 43. It is true likewise, that in God is the source not only of existences but also of essences, in so far as they are real, that is of all the reality there is in possibility. This is because the Understanding of God is the region of eternal truths or of the ideas on which they depend, and because without him there would be nothing real in the possibilities -- not only nothing existent, but also nothing possible. (*Monadology*)

God is perfect, on this view, and (as such) is a necessary being. But His creation could take different forms, since one of the perfections of God must be the freedom of Will. So, God is free to choose from among the various possible series of events (i.e. possible worlds) what will become actual, and what will remain a mere possibility:

> [T]he possible decree which is involved in the notion of the series and the things which enter into the series, and which God decides to render actual, is one thing; but the decree by which he decides to render actual that possible decree is another. (*Necessary and Contingent Truths*, p 105)

> 53. Now, as there is an infinite number of possible universes in the ideas of God, and as only one can exist, there must be a sufficient reason for God's choice, determining him to one rather than to another. (*Monadology*)

Clearly, Leibniz's understanding of contingent truths is cast in terms of possible states of affairs. God can render actual one or another world; but God first conceives of all the possibilities, and then selects one to render as actual.[1] That selected world is this one.

It is important to notice, however, that Leibniz does not understand necessity in terms of the possible worlds. From §43 of the *Monadology*, above, we can see that the foundation for the eternal truths (which are the necessary ones) is the Understanding of God. What guarantees the foundation of the modal analysis, then, is God's mind. For the eternal truths are something like the psychology of God:

> 46. We must not, however, imagine, as some do, that because the eternal truths are dependent on God, they are therefore arbitrary and depend on his will, as Descartes, and after him M. Poiret, seem to have thought. This is true only of contingent truths, whose principle is *fitness* or the choice of *the best*; whereas necessary truths depend solely on his understanding, of which they are the internal object. (*Monadology*)

We must also notice that Leibniz does not take the (non-actual) worlds to be real in any sense. They are merely conceptions in the mind of

---

[1]This is an oversimplification, as there is a difficulty in supposing that there is a difference between God willing (that is, God conceiving) and God doing.

God. Only one of them gets to be real: the actual one, which is this one. Nevertheless, it is obvious that, in Leibniz, we have the seed of an idea for understanding modal sentences. It is not far to go before we understand sentences of the form 'possibly-ϕ' as meaning that, in some world, 'ϕ' is true.

1.3 Logic, semantics and worlds

One of the difficulties for any logic which would accommodate modality was that, until the 1950's, there was nothing like a formal semantics for modal logic. During the first half of the twentieth century, there was work in the formalisation of modal systems.[2] But these efforts showed that one could construct several (non-equivalent) theories of modal inference. For example, one could have one's choice of propositional modal systems, some weaker (e.g. the Brouwer system), some stronger (e.g. S5).[3] These various systems entailed different theorems; so, different formulae would be derivable under the different systems. In the absence of any formal semantics for modal logic, logicians could not define what would qualify as a valid formula of a modal system. That meant that nobody could provide a completeness proof

---

[2]See, e.g., C.I. Lewis and C. Langford's *Symbolic Logic*.

[3]The "weakness" and "strength" of the systems is related to the assumptions they take as primitive. For our purposes, it is not really important to explore the details of the logical systems. What is important is to understand that the various systems entailed different results, which meant that one had no clear reasons to accept one system over another; but that tended to undermine the plausibility of *any* system, since none could make a compelling claim on one's intuitions.

for any modal system; it also meant that there was no obvious reason to prefer one system over another. It was not at all clear what a choice of one or another modal system would involve. That tended to undermine the (general) credibility of modal logic.

In 1940's, Rudolf Carnap[4] suggested that we could understand modal logic in terms of "state descriptions". These state descriptions, he thought, are maximally consistent sets of atomic sentences. Using state descriptions, we can understand what it means for some sentence to be necessarily true: 'necessarily-ϕ' is true just in case 'ϕ' is true under every state description. We can see how such an approach is similar to Leibniz's understanding of possibility: it considers some state of affairs as important in understanding a modal claim. Leibniz uses possible worlds to understand possibility. Carnap uses (all) state descriptions to understand necessity. We might say, then, that the state descriptions are descriptions of possible worlds; they are, in that sense, a version of a possible-worlds understanding of modality. But it is important to see, as well, that the state descriptions are different from Leibniz's worlds in an important respect. Leibniz thinks that the necessary truths are guaranteed by something outside the worlds: God. For Carnap, it is the possible worlds that *define* what necessity is: something is necessarily true just in case it is true under every state description. It is in Carnap, then, that we find the

---

[4]See, e.g., *Meaning and Necessity*.

true origin of the possible worlds which we find in later discussions of modality. For Leibniz's worlds are used to allow us to understand how God could have created the world differently. But Carnap's state descriptions (and later theorists' worlds) offer not only an explanation of possibility and necessity; they also offer a definition of those terms.

In the 1950's, several authors[5] discovered a way of interpreting modal logical operators as restricted quantifiers which range over entities to be regarded as possible worlds. That is, they offered a semantics for modal logic. The formal development of the semantics is not central to our concern[6]; but we should consider it in outline, in order to see why it offered to modal logic a "respectability" that had been missing before.

The semantics is obtained via a model structure. A model structure for a modal language, *M*, "is an ordered triple (G,K,*R*) where K

---

[5]The discussions usually mentioned are Hintikka ("Quantifiers in Deontic Logic"), Kanger (*Provability in Logic*), Kripke ("A Completeness Theorem in Modal Logic") and Montague ("Logical Necessity, Physical Necessity, Ethics, and Quantifiers"). The outline I offer is based upon Kripke's discussion in "Semantical Considerations on Modal Logic"; but the alternative approaches are not too dissimilar. There is also a helpful discussion of the history of modal logic and these semantics by Loux in *The Possible and the Actual*. The first two chapters of Forbes's *The Metaphysics of Modality* can serve as a fairly complete (albeit dry) introduction to modal logic and its semantics. D. Lewis outlines the semantical analysis of modal logic on pp 17-20 of *On the Plurality of Worlds* (henceforth, *PW*).

[6]Readers who are interested in the development of the formal semantics are directed especially to Kripke's "Semantical Considerations on Modal Logic".

is a set, *R* is a reflexive relation on K, and G ε K." ("Semantical Considerations on Modal Logic", p 64; henceforth, "SC") Intuitively we are to understand K as the set of all "possible worlds" (Kripke uses scare-quotes around the term), and G as the "real world". *R* is a relation of relative possibility, so that HRH' means that every proposition which is true in H' is possible at H. That means that the relation *R* is reflexive: anything which is *true* at H is also *possible* at H (that is, HRH).

Once we have our model structure, we develop a model which assigns, to each atomic formula (*P*) of *M* a truth-value (T or F) in each world H ε K. Formally, the model is a binary function from the atomic formulae and the possible worlds to the truth values. In other words, to each atomic formula the model assigns a truth value in every world. Once we have our model, we can determine the truth value of any formula of our language *M* by invoking some minimal rules. The rules are these (they use a standard logical notation):

(a) ~*A* is true at H if and only if *A* is false at H.

(b) (*A* v *B*) is true at H if and only if *A* is true at H or *B* is true at H.

(c) ◊*A* is true at H if and only if there is at least one world, H' where HRH' and *A* is true at H'.

(d) □*A* is true at H if and only if, for every world H' such that HRH', *A* is true at H'. (cf. "SC", pp 64-65)

Moreover, we can extend the system to include quantified modal logic. [The details of this are not important here; but in brief, the way to do so is to add modal operators to the standard predicate calculus using

only *closed* formulae, so that "assertion of $A(x)$ with free x can always be replaced with assertion of $(x)A(x)$." ("SC", p 69)] Thus do we get a fully-developed logical system. Moreover, thanks to the semantics, we can offer a completeness proof for the system. That way, we know what our acceptance of any given modal system will entail.

The effect of all of this is to enable us to understand modal sentences as sentences which quantify across worlds.[7] So, if we want to understand the sentence, "possibly-ϕ", we interpret it as saying, "There is a world such that, at that world, 'ϕ' is true." Similarly, in interpreting "necessarily-ϕ", we interpret the sentence as, "At every possible world, 'ϕ' is true." If we want, we can replace the usual logical operators (the well-known box and diamond) with the simple existential and universal quantifiers ("$\exists$" and "$\forall$"); or, we can simply read the modal operators as indicating a quantification across the worlds. What we get from the semantics, then, is a way of understanding what the basic assertions of a modal theory are about.

The details of the formal semantics are not that important for the purposes of our discussion. What *is* important is the way those semantics rely upon the notion of possible worlds. The possible-worlds semantics, we noted above, uses a model structure which incorporates the notion of possible worlds (set K). So the notion of possible worlds is

---

[7]This is, strictly speaking, an overwhelming simplification of the work that transpired between the original development of the semantics and David Lewis's "Counterpart Theory and Quantified Modal Logic"; but it captures the force of that work.

"built in" to the structure of the semantics. But so far, there are no metaphysical or ontological implications of these worlds: Kripke, for example, puts quotation marks around the term "possible worlds" ("SC", p 64), and Hintikka does the same ("The Modes of Modality", p 67). But we originally wanted to understand the notion of possibility. It is hard to believe we can clarify a concept by replacing it with references to something mysterious; yet that is all we have so far. The original, problematic concept (possibility) has given way to the concept of an object (a set of possible worlds) with which we have no contact, and for which we have no metaphysics. Such metaphysical questions become more obvious if we consider some objections to the talk of possible worlds.

1.4 Trans-world reference

We saw, above, that a sentence of the form "possibly-ϕ" is true at a world just in case the sentence "ϕ" is true at another, "accessible" possible world. This raises a sticky problem, however: what does it mean to say that "ϕ" is true at another world? There seem to be two problems here. The first is a difficulty about whether we can make any sense of modal sentences. The second is, perhaps, more serious: what is the nature of a possible world? We will deal with the first of these problems in this section; the second problem we will confront later.

The question of whether we can make sense of modal sentences itself breaks into two parts. The first issue is usually framed in terms of the "*de re* -- *de dicto* distinction". The second issue is the

problem of trans-world identity. Nevertheless, it is clear that both of these issues are related to the larger question, "Does the analysis of modality in terms of worlds make modality any more transparent?"

We can most easily understand the distinction between *de re* and *de dicto* with an example. Consider:

It is possible that the number of planets is odd.

If we interpret this as *de dicto*, we understand it to mean that "The number of planets is odd," is possibly true. If, on the other hand, we interpret the sentence as *de re*, we understand it to say *of the number of planets*, whatever that number is, it is possibly odd. This would seem to create no difficulty, except for a problem arising from identity. For we have a principle of substitutivity, which says that "given a true statement of identity, one of its two terms may be substituted for the other in any true statement and the result will be true." (Quine, "Reference and Modality", p 139; henceforth "RM") Now, consider the following:

(1) 9 is necessarily greater than 7.

but

(2) The number of planets is possibly less than 7.

(3) The number of planets is 9.

Therefore,

(4) 9 is possibly less than 7. (cf. "RM" pp 143-144).

Obviously, this will not do. For while '9' ought to be able to be substituted for 'the number of planets', because of (3), it is obvious

that such a substitution leads to a falsehood. The answer is to understand these as true *de dicto*, even though they are false *de re*. Or, as Quine says, the terms '9' and 'the number of planets' "occur irreferentially" ("RM", p 144) in (1) and (2) above.

So far, so good: we have managed to avoid confusion due to the principle of substitutivity. But now, suppose we want to replace singular terms with variables of quantification. That leads us to state the following, from (1):

(5) $(\exists x)(x$ is necessarily greater than 7).

What is the number that is necessarily greater than 7? From (1), it was 9, which is the number of planets; but now we are back to the falsehood which we avoided, above, by interpreting (1) as *de dicto* true. And, as Quine says,

> In a word, to be necessarily greater than 7 is not a trait of a number, but depends on the manner of referring to the number... Being necessarily or possibly thus and so is in general not a trait of the object concerned, but depends on the manner of referring to the object. ("RM", p 148)

But, now, given an understanding of ordinary quantification, and an understanding of modality[8], we do not have an automatic analysis of quantified modal sentences like (5). In other words, "*Necessary* greaterness than 7 makes no sense as applied to a *number x*; necessity

---

[8]When Quine wrote the original paper, he was relying on an uncritical acceptance of the notion of analyticity; but given the way he has construed the modal operator here, any analysis of modality will do. See below.

attaches only to the connection between '$x>7$'...," and some particular method of specifying x ("RM", p149).

Quine suggested that the only way out of the bind was to adopt "Aristotelian essentialism". The difficulty with essences is that they are mysterious. If we are to understand modal sentences in that way, we will have to claim that any object has some of its traits necessarily and others contingently; but the contingent traits will follow as analytically from some ways of specifying the object as the necessary traits follow from other ways of specifying it (cf. "RM", p 155). This seems intolerable.

One way around the difficulty is to reconsider how we interpret sentences like (5). Graeme Forbes (cf. *The Metaphysics of Modality*, pp 50 ff.) argues that part of the problem lies in treating necessity as a disguised metalinguistic predicate. For example, following Quine, he says that we should reinterpret

> $(\exists x)$("x is made of matter" is necessarily true)

to consist of a redundant string of symbols, '$(\exists x)$', followed by

> "x is made of matter" is necessarily true. (*Metaphysics of Modality*, pp 51-52)

Now the latter is obviously false, since 'is necessarily true' can only apply to meaningful sentences; and "x is made of matter" is no such thing. But, he says, that ignores a parallel construction which seems perfectly acceptable:

> (6) Everything is always made of matter

cannot be rewritten as

(7) $(\forall x)$("x is made of matter" is always true).

For, he says, (6) "makes good sense: it is true at a time t iff everything existing at t is made of matter at all times..." (*Metaphysics of Modality*, p 52) But (7) is badly formed: it is a string of redundant symbols, $(\forall x)$, followed by a meaningless sentence, "'x is made of matter' is always true". Since we can make sense of (6), we should reject (7) as a bad translation. Because this case is formally analogous to the case of modal operators, we should treat modal operators just as we treat trans-temporal operators. So, if we claim there can be no *de re* modal operators, we must claim that there can be no *de re* temporal operators; alternatively, we can accept *de re* modality. Otherwise, we would need to believe in trans-temporal identity, but nevertheless doubt that there can be trans-world identity.

The latter response is natural enough: we have at least an idea of what we mean by identity through time. It seems that there is some way of specifying an object such that it persists through time. The same is not obviously true of trans-world identity. I mentioned at the beginning of this section that the question of whether we can understand modal sentences could be separated into two parts; but now we see why the two parts are really forms of the same question. For if the solution to the troubles surrounding *de re* modality is to be trans-world identity, then we need to ask in what such identity consists. We seem to be forced to only one conclusion: it consists in some kind of essence of the object.

In the next chapter (§ 2.3) we shall consider David Lewis's theory of counterparts. For now, we can mention briefly that Lewis rises to Quine's charge of "Aristotelian essentialism" by pleading guilty. Lewis claims that any particular object is *identical* only to itself: it is in its own world (more on the "isolation" of worlds below, and in chapter 2). Modal sentences about that object, however, are made true by its *counterparts* in other worlds. The object's counterparts are the things which have the attribute which is the object's essence.

If we want to interpret modal sentences as true-in-a-world, then, we will be stuck with some kind of trans-world identity; or, at least, Lewis's counterpart theory. This seems to leave us with some form of essentialism. Perhaps, however, Lewis can offer a convincing story about those essences. Yet we have not clarified our notion of truth-at-a-world. Rather, we have something which looks to be even more troublesome: we seem, now, to be referring to objects which are not in our world. We have not yet begun to understand what such reference entails. What is the metaphysical status of the worlds? Perhaps a clearer description of what a world must be will clarify what we should think of the ontological status of the worlds.

1.5 The nature of worlds.

When the notion of a possible world was introduced for the purposes of semantics, we avoided too much worry about what the worlds were by using quotation marks: we pretended that the worlds carried no philosophical weight. But, as should be obvious from the discussion in

the previous section, it is difficult to offer an analysis in terms of some object without giving that object some weight. It is time, then, to consider what, at a minimum, the worlds must be like if we are to use them in our analysis of modality.

To begin with, if the worlds are to be helpful, the actual world must be one of them. We saw as much in Kripke's discussion (see §1.3), where he pointed out that among the members of set K is the actual world. But what is the actual world? When we first encountered the notion of a possible world, I suggested that we could use it as a kind of "shorthand" for possible states of affairs. For what we were trying to capture was just the notion that things could have been otherwise than they are. So, the *actual* world is the state of affairs as it actually obtains.[9] And, of course, it is *absolutely everything* that actually obtains. For now, we can say that the actual world is so complete that any statement which is actually true is so because of how things are in this world. Obviously, that means that this world contains every actual physical object, from the tiniest particle to the largest, most far-away star system. So, the actual world is "world enough": there is nothing actually true which is not true in the actual world[10]. And, as it is with the actual world, so it is with every

---

[9]If this seems circular, it is intentionally so. Lewis has a particular definition of actuality, which we will encounter in chapter 2.

[10]This is related to the fullness of logical space. There is nothing possible, on Lewis's view, which exists outside of logical

world. Each world is a complete whole, in that anything which is true at that world (and, hence, possible at other worlds) is true in virtue of the state of affairs at that world. That does not mean, of course, that everything exists at every world. It just means that there is nothing, other than the worlds, which makes sentences true.

According to Lewis, the worlds are also *isolated*, one from every other. Notice that this description of a world may not be true on every theory of worlds: Alvin Plantinga's view seems to entail that individual objects exist in more than one world. (See "Transworld Identity or Worldbound Individuals" and *The Nature of Necessity*.) But, given Lewis's solution to the problem of trans-world identity, it is a requirement that nothing exist in two worlds at once. Furthermore, there is not a possible world of possible worlds: the worlds are complete wholes in the way explained above. That means that there cannot be something outside a world *w* which makes a non-modal statement at *w* true. If there were a world of possible worlds, that would mean that the original worlds were not proper wholes, since there would be something external to them to make (non-modal) sentences true.

One more thing seems necessary to make the possible worlds semantics as useful as it initially seemed. We need to suppose that there are enough worlds for all the possibilities. Kripke captured this

---

space. Similarly, there is nothing actual which does not exist in the actual world.

notion by calling K the set of "*all* 'possible worlds'" ("SC", p 64, emphasis added); and it seems obvious enough. It would not do if our possible worlds did not include the worlds with, say, unicorns. For we are trying to explain why English sentences like, "Unicorns might have existed" are true.

Our purpose in sketching a description of the worlds was to discover whether such a description would make the metaphysics of the worlds any clearer. At the beginning of section 1.4 I noted that we need to be able to say something about the nature of the worlds, if they are to serve any explanatory role in our theory of modality. Yet, our sketch of the nature of the worlds has not helped. All we have, so far, is an indication that the worlds are things that make sentences true; and, that they are all the things which make sentences true.

David Lewis claims, in effect, that such is all we can expect. If a questioner wants to know what sort of thing possible worlds are, Lewis says,

> I can only ask him to admit that he knows what sort of thing our actual world is, and then explain that other worlds are more things of *that* sort, differing not in kind but only in what goes on at them. (*Counterfactuals*, p 85)

That response may seem extreme; but it is, in its most concise form, David Lewis's theory of modal realism. Do you want to know why a possible world makes some modal sentence true? Lewis can say it makes it true for the same reason the actual world makes statements of fact true. Why does the possible-worlds semantics for modal logic work well? Lewis can say it does because of the same reasons other semantics work

for other logics. All we need to do, Lewis will say, is make sure we are talking at the right level: existential and universal quantification for this world, and trans-world quantification for possible worlds.

"Yes," one might like to say, "but what *are* the worlds?" Lewis will respond that he already answered the question -- or, answered it as much as he can. The worlds are just the same kind of thing as the actual world, whatever that is. There is nothing more to the matter: worlds are worlds. If that seems completely implausible, Lewis says, then try to do better.

There is one more important detail to notice. Originally, we introduced possible worlds as a device to understand the notion of possibility. But, we had to explain what we needed to believe to have that explanatory device; and, Lewis says, it commits us to believing that *there are* such worlds. So, we have proceeded from trying to understand a fairly commonplace part of language -- talk of possibilities -- to vastly expanding our ontology. Of course, we might take this to mean that we have discovered how to make explicit that which was already contained (implicitly) in the phenomena to be explained -- modal language. Certainly, Lewis thinks that, if we want to have the resources that modal language offers, it forces us to accept the multiplicity of worlds.

1.6 Buying worlds

Lewis wants to sell us a philosophers' paradise. We can have any modal language we like in it. We will have no trouble with modal

idioms. We will encounter no difficulties with trans-world identity. Into the bargain, he will throw a theory of properties, and an explanation of verisimilitude. All we need to do is pay the price. That price is the acceptance of worlds. Lewis says we cannot "gain title" (*PW*, p 4) to the modal talk -- the paradise of *possibilia* -- unless we accept such talk as the literal truth. He says we cannot have paradise more cheaply. The question we must answer is whether the price is too high.

Chapter 2

Lewis's Modal Realism

David Lewis believes in worlds: he is an absolute modal realist. He thinks that the worlds exist independently of, and are of a kind with, the actual world. His argument for this position is, in itself, a fairly simple one. He argues, first, that modal idioms can be well-accommodated by using the notion of worlds, and understanding the idioms as existential quantifications across those worlds; he calls this the philosophers' paradise. Then he argues that modal realism works best to explain how the paradise is obtainable. Moreover, he argues that other approaches do not allow us the freedom of the philosophers' paradise: we cannot have paradise on the cheap. So, he says, we should be modal realists.

The complications in the argument turn up when we examine Lewis's arguments to the effect that modal realism works, and its alternatives do not. In order to understand Lewis's view, however, we must consider whether we should take modal realism seriously. So, we will have to examine just what Lewis says, and what his view entails.

2.1: Surely, you're joking?

Lewis claims that, just as the realm of sets is a paradise for mathematicians, "logical space" is a paradise for philosophers. That is why we ought to believe in the reality of worlds:

> We have only to believe in the vast realm of *possibilia*, and there we find what we need to advance our endeavours... If we want the theoretical benefits that talk of *possibilia* brings, the most straightforward way to gain honest title to them is to accept such talk as the literal truth... Modal realism is fruitful; that gives us good reason to believe that it is true. (*PW*, p4)

But Lewis is willing to concede the possibility that modal realism's fruitfulness does not give us *enough* reason to think it is true. For it could be that other beliefs we have are inconsistent with modal realism; and we might be unwilling to give up those other beliefs. Or, it could be that the same benefits as are provided by modal realism can be had at a "cheaper price": we might not need to be modal realists in order to live in the philosophers' paradise. If either of these conditions hold, Lewis's modal realism is in trouble.

Certainly, on the face of it, Lewis's theory is implausible. Lewis admits that it is very strange to believe that other worlds really exist, in just the way the actual world does:

> Modal realism *does* disagree, to an extreme extent, with firm common sense opinion about what there is... When modal realism tells you -- as it does -- that there are uncountable infinities of donkeys and protons and puddles and stars, and of planets very like Earth, and of cities very like Melbourne, and of people very like yourself, ... small wonder if you are reluctant to believe it. And if entry into philosophers' paradise requires that you do believe it, small wonder if you find the price too high. (*PW*, p 133)

That is, of course, the "incredulous stare" reply to modal realism: "It's just *too* weird." But, of course, it is pretty strange to assert that apparently solid surfaces are full of tiny, invisible holes, or that there is an invisible force which holds everything in the universe together. But, the latter two are perfectly respectable, "scientific"

beliefs. Perhaps, then, we should be modal realists; at least, we should be, if the theory is better than others we can offer to explain what we mean when we say, "Things could have been otherwise."

2.2: Worlds and modality[1]

As we saw in chapter 1, possible worlds are a useful device for understanding modal claims. That means that if 'possibly-ϕ' is true, it is true because 'ϕ' is true at some world. Modal claims are, then, just existential quantifications across worlds. In turn, when we say that something is true 'at a world, W', we are just restricting the scope of the statement, rather like what we do when we say, for example (when moving an apartment's furnishings), "All the books are in these boxes." There is an implicit restriction in the latter statement. It is not saying that all the books which *exist* are in the boxes, but that all the books which *are to be moved* are in the boxes (cf. Lewis's discussion of "in Australia" and "at W", *PW*, pp 5-6). What we have, then, in the case of "'ϕ' is true at W" is an instance of a restricting modifier.[2]

If we accept the above, we can see how easily modal language can be understood as quantification across worlds, without that quantification doing injustice to our usual intuitions about what statements mean. For example, if I look out my window and see a bare

---

[1]Compare what follows to *PW*, §1.2

[2]We should keep in mind that the restricting modifiers are to be interpreted "so as to be sensible" (*PW*, p 6).

patch on the lawn, I might make an observation: "There could be a tree there." According to the above, that means, "There is a world where there is a tree in that spot in the yard."[3] Perhaps, however, my landlord is in the room when I make the observation, and tells me that the natural gas line goes right underneath that very spot; he might say, "No, there couldn't be a tree there. That's where the gas line runs." Again, if we use worlds to make the restriction apparent, we get something like the following: "It is not the case that there is a world where trees are similar to the trees in this world, and gas lines are similar to the ones in this world, and there is a tree which shares the space in the yard with the gas line." In other words, my landlord restricts the scope: he points out that we need a world where there is both the open space in the yard *and* no gas line in that spot. By restricting the scope, he makes false what had been a true sentence: "There could be a tree there" was true for worlds where there was no gas line. Once those worlds are excluded, the statement is no longer true.

This use of restrictive modifiers also explains why Lewis rejects the notion of "impossible" worlds, where one speaks truly by contradicting oneself.[4] At such worlds, 'ϕ and not-ϕ' is supposed to be true. But if 'at W' is a restricting modifier just like 'in the

---

[3]We will see, below, Lewis's answer to the vexing problem of identity across worlds.

[4]For what follows, cf. *PW*, p 7n. We shall revisit the impossible-worlds question in chapters 3 and 5.

boxes' or 'in Hamilton' or 'around here', then we cannot speak truly by contradicting ourselves. For, in the same way that 'ϕ and not-ϕ' is false 'in the box' (and anywhere else), it is false 'at W'.

None of this is, as yet, reason to suppose that worlds *really* exist: we can understand all of the above by simply understanding the usual formal logical notation ('$\Diamond$' for 'it is possible that' and '$\Box$' for 'it is necessary that') as "interpretable as" quantification over worlds. In other words, we could use worlds as meta-logical interpretation devices without committing ourselves to any ontology at all. In the case of an unactualised possible thing -- the missing tree in my back yard -- we do not really need to be modal realists.

But what of cases that are slightly more difficult? How will we deal with a case like, "A red thing could resemble an orange thing more closely than a red thing could resemble a blue thing." (*PW*, p 13) That seems to make a comparison between things which may be parts of different worlds; as Lewis analyses it,

> For some x and y (x is red and y is orange and for all u and v (if u is red and v is blue, then x resembles y more than u resembles v)).

We can understand that fairly easily if we take the "could" in the original statement to unrestrict the comparison. So, not only is it the case that any red thing is more like an orange thing than a blue thing, but it is the case that something which is red (no matter what world) is more like anything orange (no matter what world) than anything blue (in any world). It involves a comparison of the red thing with other

things, including those which are not actual. That seems to suggest, however, that we need other-worldly things to quantify over; so, we need to believe in real worlds, populated by real things.

Worlds are similarly helpful in explaining notions like supervenience. Lewis supposes, for example, that all physical laws may be nothing more than patterns of regularity in the point-by-point distribution of properties. That would mean that two worlds could not differ in their laws without (somewhere) differing in the local distribution of properties. Now, of course, that seems very unhelpful, since it appears to say that there is no world where two worlds differ locally without differing in their laws; but there is no world where there are two worlds at all (recall that worlds are supposed to be complete and isolated; see chapter 1 and below, section 2.7). But just as with the modalised comparisons, the point of the "could" in the original locution was just to unrestrict quantifiers; the comparison is being made between worlds:

> Among all the worlds, or among all the things in all the worlds (or less than all, in case there is some restriction), there is no difference of the one sort without difference of another sort. Whether the things that differ are part of the same world is neither here nor there. (*PW*, p 17)

The things do not have to be part of the same world, since the "could" unrestricts quantifiers which would otherwise range only over actual things. That suggests, again, that there better be other-worldly things over which a quantifier might range. We need to believe in many real worlds, filled with many real things.

2.3 Who am I when I'm not at home?

One of the most common ways that modal idioms appear in our language is in the suggestion that some actual thing could be slightly different. Consider, for example, the following examples:

(i) I could have become a lawyer.

(ii) I might go to the store later.

(iii) Had I been Winston Churchill, I would never have escaped from Pretoria.

I will return to these examples from time to time, as I think they illustrate the breadth of use of modal idioms. Obviously, (iii) is the most complicated, so it will come under some scrutiny. But even (i) causes us some problems. For, if we understand it as quantifying across worlds, it means something like the following: "There is some world where I am a lawyer." The trouble is the 'I' in the sentence; since I (actually) am *not* a lawyer, in what way does it make sense to claim that is 'me' in the other world? The problem becomes more acute when we remember that, for Lewis, worlds are completely *isolated*: they are complete systems unto themselves, with nothing "other-worldly" in them (see ch 1 and section 2.7, below). It cannot be, then, that there is some super-worldly 'I' that exists across all the worlds. It must be some particular, world-bound individual in the other world who is the lawyer, *and* who is me.[5]

---

[5]This is akin to Quine's (*Word and Object* pp 245 ff) objection that the identity of unactualised possibilities is unclear.

Lewis solves the sticky problem of identity across worlds by dispensing with it. On Lewis's view, identity (strictly understood) is bound to a world. Who is it that is the lawyer in the other world, then? It is my counterpart. My counterpart in a world is the thing in that world that is most like me[6]: "[Y]our counterparts are men you *would have been*, had the world been otherwise." ("Counterpart Theory and Quantified Modal Logic", Lewis's italics, p 28; henceforth, "Counterpart Theory". Lewis credits this way of speaking to L. Sprague de Camp.) The difficulty, naturally, is that the counterpart relation is one of similarity; so while in some cases it will be easy to know when we have the relation, in many cases it will be entirely obscure. Nevertheless, we can state with some clarity some of the properties of counterparts[7].

Counterpart relations are not transitive. Suppose that *y* is an object in the actual world. There may be another world, *u*, in which there is an object, *y'*, which resembles *y* very closely (more closely, also, than anything else in *u*). Suppose, too, that there is another world, *v*, in which there is an object, *y"*, which very closely resembles *y'* (and, resembles *y'* more closely than does anything else in *v*). It might yet turn out that *y"* does not resemble *y* closely; or, it might turn out that there is something in *v* which more closely resembles *y*

---

[6]Naturally, we need to qualify this claim: some objects do not have a counterpart in some worlds. See below.

[7]For what follows, see "Counterpart Theory".

than *y*" resembles *y*. So, *y*" is not the counterpart of *y*, even though it is the counterpart of the counterpart of *y*.

Counterpart relations are not symmetric. It could be that there are two things, *x* and *y*, in this world, which share a common counterpart, *x*' (at *u*); but, the counterpart more closely resembles only one of those things (say, *x*). So, while *x*' is the counterpart of both *x* and *y*, only *x* is the counterpart of *x*'.

A thing can have more than one counterpart at a world. Two things can also share an other-worldly counterpart. But worlds do not always have counterparts for things in other worlds: a world can have an object which has no counterpart in another world, and a world can fail to have a counterpart for something in some other world. (cf. "Counterpart Theory", pp 28-9)

That gives us an idea of what counterparts are not. The postulates that Lewis sets for counterparts are as follows:

1. Nothing is in anything except a world.
2. Nothing is in two worlds.
3. Whatever is a counterpart is in a world.
4. Whatever has a counterpart is in a world.
5. Nothing is a counterpart of anything else in its world.
6. Anything in a world is a counterpart of itself.
7. Some worlds contain all and only actual things.
8. Something is actual. ("Counterpart Theory", p 27)

Lewis also provides the logical formulations of these postulates. By 2 and 8, the world mentioned in 7 is unique; Lewis further provides an abbreviation for that world's description: @. Notably absent here, however, are clear criteria for determining whether anything is a counterpart of anything else. We discover, however ("Counterpart Theory", pp 34-35, 42-43, *PW* pp 8-13), that what determines whether one object is a counterpart of another is whether it has the same essence:

> Essence and counterpart are interdefinable. We have just defined the essence of something as the attribute it shares with all and only its counterparts; a counterpart of something is anything having the attribute which is its essence. (This is not to say that that attribute is the *counterpart's* essence, or even an essential attribute of the counterpart.) ("Counterpart Theory", p 35. Lewis's italics.)

Immediately we run into a new problem. "[T]he essences of things are settled only to the extent that the counterpart relation is, and the counterpart relation is not very settled at all." ("Counterpart Theory", p 42) Lewis's answer to the trouble is to appeal to practice: the matter gets resolved, if it can be resolved at all, according to conversational context.[8] So, "any halfway reasonable statement will tend to create a context that (partially) resolves the vagueness of the counterpart relation in such a way as to make that statement true in that context." ("Counterpart Theory", p 42)

We can see, now, the way to understand the example (i). 'I could have become a lawyer' can be understood as 'At some world, W,

---

[8]We will return to the matter of conversational context in chapter 5.

there is something which is my counterpart; and that counterpart is a lawyer.' I satisfy the statement 'I could have become a lawyer' by having a counterpart who *is* a lawyer. We might wonder, however, how we might know whether I satisfy that condition: how could we know if (i) is true? Lewis will answer that question by using the device of "plenitude". We will examine that notion below. First, we should examine some other strengths of modal realism.

2.4 Just like home, only different.[9]

Worlds can help us with an analysis of counterfactuals. Lewis argues that we can understand counterfactual conditionals as invitations to consider the state of affairs under a selected counterfactual situation; such a situation is a possible world. So, he says,

> [W]e can say that a counterfactual conditional 'If it were that A, then it would be that C' is true iff C is true at the selected A-world[10]. More generally, the conditional is true at a world W iff C is true at the A-world selected from the standpoint of W. (*PW*, p 21)

---

[9]Compare what follows with *PW*, §1.3

[10]The A-world is the world at which A is true.

Of course, the selection is important: the counterfactual situation is only interesting to us if it is relevantly different *and* relevantly similar. As Lewis says,

> A counterfactual $\phi\Box\rightarrow\psi$[11] is true at a world *i* if and only if ψ holds at certain ϕ-worlds; but certainly not all ϕ-worlds matter. *'If kangaroos had no tails, they would topple over'* is true (or false, as the case may be) at our world, quite without regard to those possible worlds where kangaroos walk around on crutches, and stay upright that way. Those worlds are too far away from ours. What is meant by the counterfactual is that, things being pretty much as they are... if kangaroos had no tails they would topple over. (*Counterfactuals*, p 9)

The worlds which are relevantly similar (while still relevantly different) are the ones which are *selected* under a counterfactual conditional.[12] So, our belief in worlds allows us to understand these counterfactuals.

The ability to analyse counterfactual conditionals would not be enough to recommend modal realism, except that we need them to understand causation. For we can only understand how some effect, E, was caused by C if we understand that, had C' been the case instead of C, then E' would have occurred rather than E. That is, the notion that E depends upon C implies that, in the world where C does not happen, E does not, either. For, as Lewis argues, causal theories are motivated

---

[11]The connective symbol reads, "If it were the case that ____, then it would be the case that...". See *Counterfactuals*.

[12]For the sake of our discussion, it makes no difference how the selection works. In *Counterfactuals*, Lewis offers a formalisation of such selection.

by imagining that normal (actual) dependencies are absent, and then imagining the way things would be (cf. *PW*, p 23). That means that counterfactuals are "built into" any causal theory. So, if possible worlds help us to understand counterfactuals, they help us to understand our causal theories.

Possible worlds also help us understand verisimilitude. When we say that something is close to the truth -- whether it be Newton's laws, or some idealisation (like a frictionless plane) that we choose for theoretical simplicity -- we say that it is true in a world which is very similar to the actual world, but is different in some small way. The closeness of that world to the actual one is why the "almost-true" statement is close to the truth about the actual world: as the other world is not far (in logical space) from this one, the truths of that world are not very far from the truths of this one. The utility of such truthlike statements cannot be denied: Newton's laws are very useful, as is the device of the frictionless plane. *Possibilia* allow us to talk of such idealisations.

2.5 *Possibilia*, thought, and language[13]

Another way that *possibilia* make paradise, according to Lewis, is their utility in analysing the contents of thought and language. Here I offer but a brief summation of his somewhat lengthy arguments;

---

[13]Compare what follows with *PW*, §1.4

nevertheless, the advantages of *possibilia* are beginning to become clear.

It is obvious that our beliefs are not always settled. When we have a "gap" in our belief, we often think that, given what we already believe, there is more than one state of affairs which could be; and, these different possibilities may be equiprobable, as far as we know. Suppose I go into a windowless office building when the sky is quite cloudy, but when it has not yet rained. I may, in such a circumstance, both regret having forgotten my umbrella, and hope that it does not rain before I get home. As I prepare to leave the building, I do not know whether it is raining. I would not be surprised if, stepping outside, I get soaking wet; neither would I be shocked if the rain had not started. Approaching the door, then, I have (at least) two *doxastically accessible* worlds: one in which the sentence "It's raining" is true, and another where the sentence is false. I believe nothing (either explicitly or implicitly) which would rule out my being in either world. For all I know, my world is the one where it is not raining.

We can understand the content of thoughts as, roughly, the class of all doxastically accessible worlds.[14] More precisely, we can understand the entire content of A's system of belief as the class of A's possible worlds: A's *doxastic alternatives* (*PW*, p 28). This gives a

---

[14]For the sake of simplicity, I am going to follow Lewis and treat knowledge as a particular kind of belief. See *PW*, p 28 ff.

kind of holism to our understanding of belief, and also explains why we can have inconsistent or impossible beliefs. The holism is plain enough: since the content of belief is described by the whole class of doxastic alternatives, it includes *everything* that is true of the individual's alternatives, even if the individual is not thinking about some particular aspect (that is how the individual has implicit knowledge). But what about inconsistent beliefs? It seems that the implicit knowledge one has might preclude such inconsistencies; since doxastic alternatives specify thought content, it appears that the individual needs to be logically omniscient. Lewis gets around such a difficulty by noting that we are, to some extent, doublethinkers; so, we can suppose that an individual has several doxastic alternatives, but also gives more or less credence to the hypothesis that he is in this or that doxastic alternative. That would mean that one might give a non-zero "doxastic rating" to some belief which, upon further analysis, turned out to be impossible. In my example above, then, I might be more likely to suppose that I was in the world where "It is raining" is true. On the other hand, I might be willing to accept the report of someone who has just come into the building to the effect that the weather was fair. Both doxastic alternatives are available to me; I tend to believe that one alternative is more likely than the other. So, I will find that I worry about not having carried my umbrella that morning, while yet looking forward to a pleasant walk home.

We can make similar use of *possibilia* to analyse language. To be correct, a systematic grammar (which covers both syntax and semantics) of a natural language would have to provide a correct description of the actual linguistic practice of the language-speakers. That means that it must be able to specify truth-conditions for most (if not all) sentences of the language. But, of course, the truth-conditions of a sentence may well depend on the occasions of utterance, since speakers take into consideration the audience, context, &c. of the utterance. Now, if the speaker's audience is to believe him, he must speak such that the audience can rely on his utterances. His utterance, then, should be true for his doxastic alternatives. Now, how are we to develop our grammar so that it can tell which speakers at which times and worlds (&c.) are in a position to utter which sentences truthfully?

Lewis offers the following (cf. *PW* pp 41 ff): First, list a finite vocabulary, and assign each element within it a "syntactic category" and a "semantic value". Then, list rules for building expressions from other expressions (i.e. for combining the vocabulary); within each rule, specify the syntactic category and semantic value of the new expression (this, as a function of the categories and values of its formative expressions). One of the syntactic values will be sentences, and the truth conditions can be expressed in terms of the semantic values.

There are three ways that the truth conditions can be determined. We might want to adopt the "external" strategy, in which we

make all the semantic values speaker-relative. That means that we must rely on *possibilia*, since different speakers exist at different worlds. In this case, we build the truth value from the semantic values of the words in the sentence for the given speaker at the given world; the semantic value of the sentence will determine whether it is true for the speaker at the world.

On the other hand, we might adopt the "internal" strategy, whereby we assign semantic values without respect to speakers or worlds. We need to make the semantic values with *possibilia* included in this case, in order that the fixed semantic values determine the truth of a sentence for a speaker at a given world. The context-dependence of the semantic values gets "built in" by using *possibilia*.

The third strategy of determining semantic value is to mix the other two approaches; Lewis calls this the "moderate external strategy" (*PW*, p 42). Irrespective of which method we use, though, it is clear that we need *possibilia* to understand the context-dependence of ordinary language use. Since the truth of much of what is communicated in a given utterance is dependent upon unstated, presumably shared beliefs, we must be ready to evaluate the truth of a sentence relative to the relevant worlds (or doxastic alternatives). *Possibilia* allow us to do that.

2.6 *Possibilia* and properties[15]

Once we believe in possible worlds and possible individuals, all we need is to believe in set-theoretic constructions of things we believe in, and we have entities suited to play the role of properties. A property is simply the set of *all* its instances, this- and other-worldly.

In this case, one cannot object that different properties may happen to be co-extensive. Accidentally co-extensive properties in one world -- or even, in many worlds -- are still not really co-extensive, since there will be some world in which they are not. Under modal realism, any way a world could be is the way some world is; so, if two properties are not necessarily co-extensive (i.e. if they are not really one property), there will be some world where they are not co-extensive.

The contingency of properties is something that we get by considering the members of the set, and their counterparts. So, for example, Nellie may have the property of being red-haired actually. She is, then, herself a member of the set of red-haired things. Her counterpart, however, may have black hair. It is contingent whether Nellie is red-haired because some of her counterparts are not among the red-haired things.

Relations work in the same way. An ordered pair of related things is just an instance of the relation; the relation is the set of

---

[15]Compare what follows with *PW*, §1.5

all such related things. These constructions do not admit of degree. So, the relation 'A is physically close to B' is better expressed as 'A is *n* distance from B'. A and B stand in the relation of being so far apart, just like everything else, both this-worldly and other-worldly, which is that far apart.

There are, of course, abundant properties: if nothing else, for any set whatever, there is the property of belonging to that set. Indeed, there is no limit to the abundance of properties. But these properties seem to be *too* abundant. So (Lewis says, *PW*, pp 59-60), we also have the conception of *sparse* properties. There are "just enough" of these "to characterise things completely and without redundancy." (*PW*, p 60) These sparse properties are a very small minority of the abundant properties; but they are not a different kind of entity. If a property belongs to the small minority, Lewis calls it a natural property. A property is natural *simpliciter*, and not in relation to another world. Some properties are *perfectly* natural: an example is mass or charge of a particle. All perfectly natural properties are intrinsic properties ("which things have in virtue of the way they themselves are", *PW*, p 61). The reverse is not true: some intrinsic properties are not perfectly natural. Now, two things are duplicates if and only if

> (1) they have exactly the same perfectly natural properties, and (2) their parts can be put into correspondence in such a way that corresponding parts have exactly the same perfectly natural properties, and stand in the same perfectly natural relations. (*PW*, p 61)

But, Lewis says, an intrinsic property can never vary between duplicates. That means that duplicates always share all their intrinsic properties. (We will have occasion to return to duplication in chapter 3.)

The ability to explain exactly what constitutes a property (or relation) is desirable, since we regularly speak of properties, and even have occasion to quantify over them. Talk of *possibilia* gives us the resources to explain properties. It is one more way, Lewis argues, that the world of *possibilia* is the philosophers' paradise.

2.7 Worlds apart

In the last chapter, we saw some of the motivations for supposing that there are possible worlds, and considered briefly how Lewis describes his worlds. There are two qualities of logical space, however, which demand emphasis. The first is the isolation of the worlds. The second is the plenitude of logical space.

Each world is of a kind with the others. But, the worlds cannot overlap; for if they did, they would not really be whole alternatives to one another. So, Lewis says,

> for any two possible individuals, if every particular part of one is spatiotemporally related to every particular part of the other that is wholly distinct from it, then the two are worldmates. (*PW*, p 70)

Moreover,

> things are worldmates iff they are spatiotemporally related. A world is unified, then, by the spatiotemporal interrelation of its parts. There are no spatiotemporal relations across the boundary between one world and another; but no matter how

> we draw a boundary within a world, there will be spatiotemporal relations across it. (*PW*, p 71)

This means that the worlds are utterly separated in logical space. Nothing exists across worlds, and anything that exists must be in one and only one world. Anything which needs to be in more than one world at a time is not a possible thing.

Worlds are nothing more than the (mereological, maximal) sum of their parts. That means that there can be no world where there is *absolutely* nothing, which makes it necessarily true that there is something. Lewis admits that such is an uncomfortable consequence of his view, but maintains that it is not that bad (*PW*, p 74). He also admits of a problem in the notion of "spatiotemporal relations" when we consider, for example, a truly Newtonian world: distances of spacetime are only one distance in our world, but two in the Newtonian world. He circumvents that problem by suggesting, wherever spatiotemporal relations fail, there is an analogical spatiotemporal relation. His solution is, admittedly, messy, but it is enough (*PW* pp 75-76).

2.8 Worlds a-plenty[16]

In order that modal realism make any sense, we need a principle of plenitude: there must be possibilities enough. The reasons for that should be obvious. For the notion of possible worlds is supposed to provide the basis for a systematic theory of ordinary uses of modal idioms:

---

[16]Compare what follows with *PW*, §1.8. The unicorn and dragon example is Lewis's.

> In trying to improve the unity and economy of our total theory by providing resources that will afford analyses, for instance of modality as quantification over worlds, I am trying to accomplish two things that somewhat conflict. I am trying to *improve* that theory, that is to change it. But I am trying to improve *that* theory, that is to leave it recognisably the same theory we had before. (*PW*, p 134)

Clearly, then, it will not do for modal realism to make false, *without explanation*, what most of us take to be true. Lewis can, for example, explain why the ordinary belief, "Only actual things exist," is false; he says that any modal sentence is a quantification over worlds, so that the objects, while not actual, nevertheless exist. It would be more troublesome if we could, in fullness, imagine a world which contained both unicorns and dragons, but which cannot exist according to modal realism. Logical space must be full; we need a principle of plenitude.

Lewis offers a principle of recombination. The principle is more or less that anything can co-exist with anything else (at least provided they occupy distinct positions in spacetime), or fail to co-exist with anything else. Now, it is not exactly that things from two different worlds can co-exist; for that would mean that there were trans-world identity relations. As we saw in section 2.3, Lewis normally handles such trans-world comparisons with counterpart relations. The difficulty with that approach, he says, is that counterparts are united primarily by extrinsic similarity; in particular, the match of origins is important (*PW*, p 88). Now, we can imagine a dragon, and we can imagine a unicorn. It is nevertheless

possible that the world of the dragon is simply too alien[17] to the world of the unicorn, such that there could not be a world where counterparts of the dragon and the unicorn could co-exist. For the unicorn's original world might be very different from the dragon's world. In that case there might be nothing in the dragon-sharing world which would match the unicorn in origins; that means the unicorn has no counterpart there: "to the extent that the counterpart relation heeds extrinsic similarities, we take [counterparts] together with their surroundings." (*PW*, p 89) Therefore, Lewis prefers to fill logical space with duplicates. (Recall that duplication is a matter of sharing all perfectly natural properties, or having all the same intrinsic properties.) A duplicate can co-exist with a duplicate of anything else, so long as all the duplicates can fit in the world[18].

Lewis's principle of plenitude seems to give us just enough. We can imagine fancifully because we can simply put together (or take apart) various bits of this world: we imagine a unicorn by imagining a horse and a horn, and then imagining that the two are attached one to the other. Naturally, we can turn out to be wrong in our imagination, as we do not imagine everything about the world in question. We can

---

[17]An alien natural property is one that "is not instantiated by any part of this world, and ... is not definable as a conjunctive or structural property build up from constituents that are all instantiated by parts of this world." (*PW*, p 91) Something can only be alien relative to some world.

[18]cf. *PW*, pp 89-90.

imagine an impossibility, and just not notice that we have done so. The recombination of duplicates prevents such impossible objects from populating logical space.

I suggested in section 2.3 that plenitude also gives us a way to know the truth of such statements as my first example, "I could have become a lawyer." Now we can see how plenitude makes such a statement true. The property of being a lawyer is just the set of instantiations of every lawyer in every world. Now, since we have as many worlds as we need, there must be some world in which my counterpart is a lawyer. Therefore, when I say, "I could have become a lawyer," I speak the truth.

2.9 Remaining worries

Since Lewis thinks that possible worlds are real, he seems to make everything actual. Lewis denies that by making a claim about what the word 'actual' means. For Lewis, to say that something is actual is to say that it is one of your worldmates. Anything which is in a world is actual at that world. Were it possible, if one were to go to another world, there would be new things which qualified as actual; but, among them would be none of one's former worldmates.

Lewis also thinks that all the worlds are concrete; but he does not think that means very much, as he seems to think that calling something abstract is little more than a way of saying "don't worry" about that thing (cf. *PW*, §1.7). We will have some cause to refer to the distinction between the concrete and abstract below; since nothing

in my later argument turns on the distinction, we can simply accept that, whatever the worlds are, and whatever 'concrete' means, the worlds are concrete.

2.10 Other possibilities

I said at the beginning that Lewis's argument for modal realism rests partly on the claim that its alternatives cannot be true. Since we want the philosophers' paradise he offers, and we cannot have it by any other means, he says it is worth the (ontological) cost to become a modal realist. We should therefore consider his arguments against what he calls ersatzism.

Lewis offers some reasons to suppose that such alternatives to modal realism will not work[19]. All of these amount to an attempt to show that "ersatzism" is false. Ersatzism is the claim that a modal idiom does not refer to some "real" thing, but instead, to some other thing, which has no independent existence: an ersatz world, which does the work that Lewis's real worlds do in the analysis of modality. So, the possible world where I am a lawyer is some "abstract" entity, rather than the "concrete" one in which Lewis believes.

---

[19]In what follows, I have attempted just to take Lewis's arguments at face value; I have not attempted to defend his opponents at all. I should nevertheless note that his characterisations of some alternative positions are, it appears, somewhat less than fair. It also seems that he simply assumes that the entities must be abstract, and not "non-existent", ones. I take that to be due to a belief in the ontological commitment of existential quantifications. See chapter 4.

The key to Lewis's objection is his view that we cannot take our modal idioms as unanalysed primitives, since that would not be a *theory* of modality, but an abstinence from theorising (*Counterfactuals*, p 85). So, for example, he considers that possibility is really just consistency disguised: 'possibly ϕ' turns out to mean that ϕ is a consistent sentence. Now the difficulty is in saying what consistency is:

> If a consistent sentence is one that could be true, or one that is not necessarily false, then the theory is circular.... If a consistent sentence is one whose denial is not a theorem of some specified deductive system, then the theory is incorrect rather than circular: no falsehood of arithmetic is possibly true, but for any deductive system you care to specify either there are falsehoods among its theorems or there is some falsehood of arithmetic whose denial is not among its theorems. If a consistent sentence is one that comes out true under some assignment of extensions to the non-logical vocabulary, then the theory is incorrect: some assignments of extensions are impossible, for example one that assigns overlapping extensions to the English terms 'pig' and 'sheep'. If a consistent sentence is one that comes out true under some possible assignment of extensions, then the theory is again circular. (*Counterfactuals*, p 85)

It will not do, then, to try to replace possible worlds with maximal consistent sets of sentences. That is, nevertheless, the approach of linguistic ersatzism. A linguistic ersatzer takes a possible world to be something like a complete, consistent novel. That raises two difficulties. As we noted above, Lewis is concerned that modality not be taken as primitive. But, if we try to analyse modality as referring to "ersatz" worlds, which are just maximally consistent sets of utterances (in a world-making language), we are making what seems to be a modal distinction: "a set of sentences is consistent iff those

sentences, as interpreted, *could* all be true together." (*PW*, p 151, Lewis's italics) So, we cannot get a definition of possibility without either falling into circularity on the one hand, or defining possibility in terms of truth in some possible world. Lewis thinks it is obvious which of these to prefer, and so he opts for the possible-worlds understanding of possibility; those worlds have to be real things if the analysis of modality is not to be circular.

There is, additionally, the difficulty that an ersatz world-making language, if it is a language that can be specified, will be limited by the specifier, which must be in the actual world. In other words, I need to be able to develop the world-making language in the actual world. That must mean that the world-making language cannot distinguish all the possibilities it should, for it will not have the vocabulary to describe the properties completely alien to the actual world, but which nevertheless apply to objects in the ersatz world.[20] Of course, if it is not a language that can be specified, then it is

---

[20]Lewis's objection really has two parts, one of which is an objection about indiscernible possible individuals, which I do not consider here in the interests of charity; for I cannot see why it is a problem at all (it appears on pp 157-58 of *PW*). The objection I consider here depends upon whether one thinks there can be a true and complete description of the world. Clearly, if one is not bothered by a language which cannot give one (and only one) true and complete description of the world, this objection will not get very far. It seems at least possible, however, that one could believe in something like ersatzism without accepting that the ersatz world has to be complete. I will consider that approach briefly in chapter 5.

difficult to see how we can define possibility in terms of that language; and, the ersatzist programme fails.

Lewis takes a similar line with "pictorial ersatzism". It would seem that, if worlds cannot be maximal consistent sets of sentences, they might nevertheless be "ersatz" by being complete, three-dimensional pictures of possible states of affairs. A picture, after all, is isomorphic with that which it represents. These are ideal pictures, so there are no properties which we want in the possible world which are not represented by the picture; that gets around the problem of the actual person being unable to say anything about alien properties. The only provision is that the pictures are "abstract" entities, rather than the "concrete" entity which is the actual world. Now, the isomorphism ought to guarantee that we do not have the sorts of incompleteness problems that linguistic ersatz worlds have. Moreover, since a picture is not made of propositions, there ought to be no difficulty with circularity. Unfortunately, Lewis shows that the approach still relies upon primitive modality (which renders it circular).[21] Consider, for example, something that does not exist: say, a talking donkey (cf. *PW*, pp 167-8). According to pictorial ersatzism, there could be a talking donkey iff there is some ersatz world in which

---

[21]Again, Lewis also offers a criticism based upon the actualisation of indiscernibles, but I do not consider it, for the same reason as above.

is pictured a talking donkey. Now, what is the talking donkey? It is not isomorphic to any actual donkey, since donkeys do not talk.

> What makes the thing an ersatz talking donkey is just that it *could* have been isomorphic to a talking donkey that was part of the concrete world, and it *would* have been if the concrete world had been different....
>
> In short, there might be a talking donkey iff there might be a talking donkey isomorphic to some part of some ersatz world. (*PW*, p168)

The last sentence is obviously no analysis whatsoever. Moreover, Lewis argues, the theory does not even offer a simpler ontology. He asks what the difference is between the actual world and the pictorial ersatz worlds. After all, the pictorial worlds must be somewhat similar, since they are supposed to be isomorphic with the actual worlds (or, even, the way the actual world would be if the ersatz world were actualised). The difference seems to be that one of them is concrete -- this one -- while the rest are abstract. The problem, now, is to say what these terms mean.

Lewis offers four possibilities for comparing concreteness and abstractness.[22] The Way of Example claims that concrete things are like donkeys and stars, while abstract things are like sets and numbers. Of course, the ersatz things are supposed to be just like concrete things, only abstract; so, this Way must be wrong.

---

[22]These are the "four ways", which appear first in *PW* pp 82-86. The objection that follows is at pp 171-174.

The Way of Conflation states that the distinction is nothing more than the distinction between individuals and sets, or particular individuals and any other thing. But it seems that pictorial ersatz worlds must be particular individuals. For two possibilities must be able to differ in their particulars, which would mean that their (ersatz) worlds must differ. So the ersatz worlds must be particular individuals if they are to represent particular, different possibilities. That means that the abstraction of pictorial ersatz worlds cannot come from their not being particular individuals; the Way of Conflation is no help.

The Negative Way maintains that abstract entities have no spatiotemporal location and do not enter in to causal interaction. Now, since the ersatz worlds are to be (at least in some cases) isomorphic with the actual one, at least the parts of these things must have spatiotemporal relations. Moreover, there must be causal relations between the parts of the ersatz worlds. But the parts of ersatz worlds must be abstract (otherwise, they are concrete, which makes them actual, not possible individuals). Therefore, these worlds are not abstract according to the Negative Way.

That leaves only the Way of Abstraction, which claims that an abstract entity is whatever is gained when one subtracts specificity

from a concrete object.[23] Now, it is apparent that specificity is not what is missing from the abstract objects; but, something is. Lewis calls it *vim* (*PW*, p 173). Some ersatz worlds are abstract, then, in that they are just like the concrete world, only without the *vim*. Other abstract entities are not really isomorphic with the concrete world, but we call them abstract because they are *vim*-less.

It is apparent that the last approach is not really ontologically simpler than modal realism. Just like modal realism, it has lots of extra entities -- albeit abstract ones -- which it entails. But, it does not offer any clear account of what *vim* is: the difference between the actual world and other possibilities is obscure, if not completely mysterious. Now, we originally wanted to have the philosophers' paradise "on the cheap": we wanted the advantages of *possibilia* without the theoretical cost. But pictorial ersatzism replaces Lewis's real worlds, numerous as they are, with at least as many things which are different in kind. So, we get out without an expansion of the number of concrete things; but we wind up with a lot of completely mysterious, abstract entities instead. The advantages of pictorial ersatz worlds are not at all clear.

Increasingly, it looks like we may have no option but to accept Lewis's strange ontology, since it is the only thing that will

---

[23]Lewis's objection will work regardless of what abstraction is: we could as easily say that an abstract object is whatever is left when one takes the concreteness away from a concrete object.

get us modality without the difficulties we see above. There is, however, a third position: that alternative, ersatz worlds are not the composite worlds of linguistic or pictorial ersatzism. On the contrary, the worlds are simples. We cannot say anything positive about their nature. We can only state that, whatever they are like, they are abstract (i.e. not concrete) entities. We must accept that modality is primitive on the theory; but, suppose we are willing to accept that. The trouble starts in defining what the modal idioms mean. For the difference, under this theory, between the actual and the non-actual is just the difference of elements which are "selected" by the concrete world as opposed to those which are not so selected. Now, selection seems to be a relation: there is some relation between the selected elements and the selecting world. That selection can either be an internal or an external relation.

In the case of an internal selection relation, the concrete world selects an element due to what goes on in the concrete world and the intrinsic nature of the selected element. If that is the case, then various elements would be necessarily selected whenever their intrinsic natures accorded with the state of affairs in the concrete world. This is not too helpful: suppose we say, "Possibly, there is a talking donkey." Then,

> [t]here is an element such that, necessarily, it is selected iff a donkey talks; that element has some distinctive intrinsic property; that property is named 'representing that a donkey talks'; the property with that name singles out the element that, necessarily, is selected iff a donkey talks. (*PW*, p 178)

In other words, an internal selection relation will be circular.

We might now like to say that the selection relation is external. This relation does not depend upon the intrinsic nature of the elements; indeed, the elements may have no distinctive, intrinsic natures. Rather, an element is selected just in case it happens to be selected.[24] Various elements make up the world; so, if there are talking donkeys, they arrange in some particular way. If I am Winston Churchill, and fail to escape from Pretoria, the elements are also arranged in some particular way. The connection between the concrete world and these elements must be a necessary connection (since it is the selected elements which constitute the way the world is). On the face of it, this is not really a theory. It uses the abstract elements to sweep any difficulties under the carpet of abstraction. In other words, when we ask, "Why didn't I get to be Winston Churchill?" the response is easy, but empty: the elements of the concrete world were not so arranged. Now, this selection of elements seems, as Lewis charges, magical. Why is *this* the actual world? If the selection relation is internal, it is because the elements which are selected (for some reason) are the ones whose intrinsic properties are just like the way things are in the actual world. If the selection relation is external,

---

[24]This is probably a stronger statement than even Lewis would accept (cf. *PW*, pp 180 ff.); but, I do not think it matters. It is entirely unclear how a relation which has nothing to do with the nature of the actual elements could somehow be constitutive of what is actually the case.

the concrete world is actual because the elements are arranged in just such a way as to make the actual world as it is. No matter what we do, magical ersatzism, as Lewis portrays it, seems to be a minefield of confusion.

2.11 Should we be modal realists?

At the beginning of the chapter, I said that Lewis's argument works two ways. First, Lewis shows that modal realism works. Clearly, he can do so. We saw, for example, that modal realism is useful for analysing modality, thought and language, properties, and counterfactual conditionals. We also found that it has the resources to deal with any troubles about trans-world identity, in the form of counterpart relations. We noted, as well, the ways it deals with difficulties like the plenitude of logical space, the isolation of worlds, and the notion of actuality.

The second part of Lewis's argument is an attempt to show that the alternatives will not work. We considered briefly his arguments against ersatzism, and saw that his arguments seem to be good ones. We seem, then, to have two choices: we can abandon *possibilia*, or we can become modal realists.

I would agree, except for a few difficulties. First, I think Lewis's view has some internal difficulties which make the view unacceptable. Second, I think it rests on a misunderstanding of how language works. Finally, I think we can come up with another approach, which will allow us all the freedom we like to talk of *possibilia*,

without committing us to any unhappy ontological view. If I am right, we really will have paradise on the cheap.

Chapter 3

Will it Work?

3.1 My counterparts, part 1

In the last chapter, I offered three examples which, I said, were representative of some important, everyday uses of possibility:

(i) I could have become a lawyer.

(ii) I might go to the store later.

(iii) Had I been Winston Churchill, I would never have escaped from Pretoria.[1]

I mentioned, there, that Lewis can easily explain how (i) can be true, by appealing to possible worlds and referring to my counterparts. We interpret (i) as saying that there is a world at which my counterpart (who is otherwise relevantly similar to me) is a lawyer.

We can see a similar way to analyse (ii): we specify a (later-than-now) point in time, and then say that, at some world(s) at that time, my counterpart goes to the store. That world, of course, might turn out to be this one; in which case, it is not my counterpart who

---

[1]These examples are not the simplest we could consider: "There might have been a talking donkey" raises no problems of trans-world identity of individuals. But those cases are going to be easy for anyone. The simplest, of course, is that the speaker simply puts together two actual concepts -- "donkey" and "something that talks" -- and comes up with a composite concept ("talking donkey"). I suggest that we can learn more about our notions of possibility by examining cases like the ones I offer.

goes, but me.[2] So, (ii) turns out to be (more or less) the same sort of case as (i), except that in (ii), we have no information about the actual world. When I claim that I could have been a lawyer, I make the claim knowing full well that I am not actually a lawyer. When I claim that I might go to the store, however, I simply do not know whether I actually *will* go to the store. To make it more explicit, then, we might try to state the time of my going to the store:

(ii') I might go to the store at 5 p.m. on October 11.

Now, at 6 p.m. on October 11, we will be in a position to know whether I do go in the actual world. Notice, however, that our lack of knowledge about the actual case here makes no difference: (ii') is true irrespective of actuality, because there is some world in which my counterpart goes to the (counterpart) store at the (counterpart) time. Even if the actual world is destroyed before October, or I (actually) get hit by a bus on October 10, (ii') is true.

But what about example (iii)? It seems to be the most difficult. How can Lewis address such an example? At first, the example seems to be amenable to the normal possible-worlds analysis: there is a world such that my counterpart is Winston Churchill, and in that world Winston Churchill fails to escape from Pretoria. When we state it baldly, however, the trouble is plain: whose counterpart is in

---

[2]This seems to commit Lewis to some form of hard determinism about individual lives. But given that "it is not contingent what conditions the entire system of worlds does or doesn't satisfy" (*PW*, p 125), Lewis must hold a view that determinism is true.

that world? Is it the counterpart of Winston Churchill? Is it my counterpart? Is it both? No matter how we answer these questions, we seem to create more difficulties than we solve.

Consider the case in which the counterpart (call him WC*) is that of the this-worldly Winston Churchill. We will have an easy time explaining how WC* fails to escape from Pretoria even though Churchill succeeds: we use the same strategy as we used for example (i). The actual world and the other world are excellent counterparts until the point when Churchill (actually) escapes from the Boers; at that point, WC* fails to escape. There is merely the normal difficulty in telling what other-worldly thing is Churchill's counterpart. In such a case, there is little problem: the histories of the two men are very similar, so they are excellent counterparts.

The problem, in case (iii), is to say how it is that *I* am the one who does not escape. For there does not seem to be a way in which a counterpart of me can also be a counterpart of Winston Churchill. (A symmetrical analysis can be made of the case wherein my counterpart fails to escape from Pretoria: how is that counterpart also a counterpart of Winston Churchill?) After all, Churchill and I share nothing whatever in our origins, except that in both cases the parents were human. Obviously, that cannot be enough similarity that we could share a counterpart, for it would mean that everyone's counterpart is also the counterpart of everyone else. Indeed, the problem seems to be that, since a counterpart of x has the essential attribute of x, it is

very unlikely that two very different actual people could share a counterpart. That is especially true given Lewis's claim (*PW*, p 88) that match of origins has "decisive weight" in uniting counterparts. It seems very unlikely, then, that Churchill and I could share a counterpart. For whatever my essence, and whatever that of Churchill, I cannot believe that "Churchill and I are essentially different" is too strong a claim. Since counterpart and essence are interdefinable, that means that Churchill and I cannot share a counterpart.

3.2 Duplicates and counterparts

If I am right, normal counterpart relations cannot explain cases like example (iii). But Lewis faced a somewhat similar case with unicorns and dragons. Recall that the unicorn's counterpart and the dragon's counterpart were not compossible because, Lewis said, the unicorn-world and dragon-world were too different. They could not both have a match of origins in the shared world. He circumvented the problem (in order to save plenitude) with the device of duplicates. Now, the duplicates have all the same intrinsic properties, whereas counterparts are united mostly by extrinsic similarity. So, perhaps the world which makes (iii) true is one where the *duplicate* of me is the *counterpart* of Churchill: the counterpart would be that thing which is extrinsically "similar enough" to Winston Churchill, but which is intrinsically the same as me. Such an approach appears to be a solution. My duplicate does not need to have the same extrinsic relations as I do, and Churchill's counterpart needs to have only enough

extrinsic similarity to be united with Churchill. Now, however, we can see what seems fishy in Lewis's attempt to use duplicates in his principle of recombination. For if the duplicate of me is not somehow "recognisably me" -- i.e. if it could not stand as a counterpart for me -- then the claim that it is "me" who is Winston Churchill's counterpart seems terribly contrived.

To make that explicit, recall that duplicates are things that share all their perfectly natural properties. These "perfectly natural properties" are terribly elusive; what are the perfectly natural properties of, say, a unicorn? Of course, the perfectly natural properties are supposed to be intrinsic properties, which are those "which things have in virtue of the way they themselves are." (*PW*, p 61). We may take such properties as primitive, or offer some theory of their constitution; neither will provide any more clarity.

It is obvious that things have some properties "in virtue of the way they themselves are." But even if we readily accept that there are intrinsic properties of things, we might yet wonder which of those properties is perfectly natural. Lewis seems to think that the perfectly natural properties are things like mass or charge of particles: "The colours, as we now know, are inferior in naturalness to such perfectly natural properties as mass or charge..." (*PW*, p 61) But that raises a difficulty: is the unicorn-world made up like this one? Do particles have mass and spin in the unicorn world? Who can tell? So, either Lewis is committed to the claim that all the worlds share the

same fundamental physical structure, or he must admit that he can tell us nothing about what the perfectly natural properties are.

It does not seem acceptable for Lewis to claim that all the worlds have the same basic structure. For it does not seem beyond the realm of possibility that some worlds have no matter: there could be, for example, a spirit-world. Perhaps Lewis can simply deny that such a world is possible; but then he needs to be able to explain why his view does not agree with "common sense" (see below, and *PW* p 134). And, in order to show that his theory captures the "real" possibilities (while "common sense" includes some false possibilities), he will have to use the notion of perfectly natural properties. But that would beg the question: the notion of the perfectly natural properties is just what is being questioned.

Perhaps Lewis will want to say that the perfectly natural properties of a unicorn might be constituted differently, but they are still the most primitive properties in that world. (That would not mean that the perfectly natural properties of the unicorn are only perfectly natural in the unicorn-world, since Lewis denies that such is possible. Rather, we could imagine that there are more perfectly natural properties than are instantiated in this world.) But if the perfectly natural properties of an other-worldly thing can be different in kind from the this-worldly perfectly natural properties, then the notion of perfectly natural properties is even more obscure than it originally seemed. That obscurity ought to come as no surprise. Lewis says (*PW*, p

60 n) that the natural properties are so named after the more familiar "natural kinds". Of course, natural kinds are notoriously elusive themselves. It is no wonder that the perfectly natural properties are mysterious.

What is worse is that the perfectly natural properties were supposed to provide us with the duplicates which were to fill logical space. We said, for example, that it was the duplicates of the dragon and unicorn which were compossible, since their counterparts could not share a possible world. That was because the extrinsic similarities which would make for counterparts of the unicorn and dragon were not compossible: the dragon-world and the unicorn-world were too different. But in what sense do we have "the unicorn" duplicated in a world with "the dragon", if the extrinsic similarities do not hold? Perhaps we would like to say that, if we had the two duplicates side-by-side, we could not tell the difference. So far, so good: we have the unicorn (or something intrinsically just like it) in a world with the dragon. But now the trouble arises. For it does not seem that all the duplicates of something can "stand for" that thing, in the way that counterparts can (more on that in a moment). But if modal claims are in no important way linked to what is true of a thing's duplicates, then the device of duplication seems to do nothing to fill logical space. In that case, Lewis's theory starts to look very different from the "common sense" theory, the credence of which he is trying to inherit (see *PW*, p 134).

It cannot be that every duplicate "stands for" all its other duplicates. We can see as much easily, by considering a case like the following: "What would my brother be like if I did not exist, but my brother had all and only my intrinsic properties?" On a possible-worlds analysis, we would understand this case as referring to a world where my duplicate is my (actual) brother's counterpart. And, in that world, I do not exist: there can be nothing which "stands for" me there. Obviously, then, there are at least some cases where duplicates cannot "stand for" each other, in the way that counterparts can.

The mistake lies in thinking that duplicates make modal sentences true. If we recall the definition of duplicates, we remember that two things are duplicates if and only if they have the same perfectly natural properties. It is from there that we conclude that intrinsic properties cannot differ between duplicates. But since the perfectly natural properties are a very small minority of all the properties, then it seems that two things do not need to be terribly similar at all in order to be duplicates. So modal statements about *x* cannot be made true by *x*'s duplicates. This raises a new problem. Recall that the point of invoking duplicates was just to fill logical space: places where extrinsic relations would, by definition, preclude counterparts get a duplicate instead. If we accept that duplicates are not united with the object of which they are duplicates, then we cannot use the duplicates to conclude anything about the duplicated objects. In other words, when the unicorn and the dragon find themselves together

in the same world, they are no longer the unicorn and the dragon. They are unicorn-like and dragon-like objects. That means that the unicorn and the dragon really cannot co-exist; and the plenitude of logical space is lost.

Perhaps plenitude is not lost. If it turns out that (logically) the counterparts of dragons and unicorns are not compossible, and we cannot use duplication to replace those counterparts in our theory, then so much the worse for our beliefs about unicorns and dragons. But Lewis is already fighting a battle with credence, so he needs to be able to account for the beliefs we had before we came to modal realism. If, before, we were able to imagine the dragon and the unicorn playing together in some enchanted forest, Lewis must be able to convince us that belief was false. The more often such cases crop up, the more difficult his time is in making his view acceptable. As he says, "A worthwhile theory must be credible, and a credible theory must be conservative. It cannot gain, and it cannot deserve, credence if it disagrees with too much of what we thought before." (*PW*, p 134)

In fairness, however, Lewis's view probably does not need duplicates very badly. His example of the unicorn and the dragon was designed so that the objects' counterparts were not, *ex hypothesi*, compossible. It does not seem unreasonable to say that things which cannot exist together cannot exist together: logical plenitude is not forsaken that way. So perhaps he could abandon the notion of duplication altogether, thereby saving himself the trouble of defending

the mysterious perfectly natural properties. That strategy will cause a new trouble, however; so he is stuck with duplicates.

To see why, consider how logical plenitude would be satisfied without duplicates. We could say that a counterpart of anything can co-exist with a counterpart of anything else. But Lewis introduced the conflicting unicorn- and dragon-worlds in order to show that there seem to be cases where counterparts *cannot* co-exist. We could, of course, get around that by subjecting the co-existence of counterparts to the strictures of logic. That makes the theory circular: a counterpart of anything can possibly co-exist with a counterpart of anything else just in case it is logically possible that the counterparts co-exist. So, counterparts cannot fill logical space.

The way out of the bind is obviously with duplicates. A duplicate of anything can co-exist with a duplicate of anything else. However, the duplicate need not "stand for" that of which it is a duplicate. Let us accept that duplicates of *x* do not make modal sentences about *x* true. Then, if *x* has a duplicate with property P, but no counterpart with P, then 'possibly, P*x*' is false. Let us suppose, for example, that Lewis's unicorn-duplicates and dragon-duplicates can be combined in a world. Now, if we suppose that these are duplicates of the only ways unicorns can be, and the only ways dragons can be, then we have the combination of the duplicates of the only possible dragons, and the only possible unicorns. But, in spite of the combination of the duplicate-unicorn and duplicate-dragon, the statement, "Possibly,

unicorns and dragons co-exist," *is false*. For if only counterpart relations make modal sentences true, and there is no world where there is both a *counterpart*-unicorn and a *counterpart*-dragon, then no sentence about such a world will be true.

So, we use duplication to fill logical space: a duplicate of anything can co-exist with a duplicate of anything else. But even if we suppose that there are only two worlds, a duplicate, *y*, of object *x* may yet be a counterpart of some other, quite different object, *z*. Or, it may be a counterpart of nothing at all. And, since counterpart relations are determined primarily by extrinsic similarity, it appears that the matter is to be settled by empirical investigation, even if that investigation is only a logical one. We will return to that matter below, in section 3.4

Notice, however, that we now have preserved the principle of logical plenitude, according to which there are enough worlds to account for all possibilities, by invoking a perfectly obscure notion: that of perfectly natural properties. Lewis appears to be agnostic about what these things are. It appears that he cannot afford to be so sanguine.

3.3 My counterparts, part 2

We originally raised the issue of duplication because of the problems in offering a modal-realist interpretation of statements like (iii), where I am Winston Churchill, and fail to escape from Pretoria. We appear to be no closer to our goal. It seems likely that the original analysis I proposed was a bad one. We might now wish to

disavow the example as meaningless; after all, our difficulty in analysing the notion of 'my being Winston Churchill' is what caused us the trouble above. But (iii) is an example of a construction we use all the time in regular speech. We need to be able to say something about it, then, or our theory will leave out a significant portion of the phenomena to be explained. Yet there is an analysis of the example which skirts some of the problems. We could simply say that (iii) can be analysed as saying, "Were I faced with the situation that Winston Churchill faced, I never would have escaped from Pretoria." This seems like a good analysis: it is similar to the oft-heard, "If I were in your shoes..." The possible-worlds analysis, then, would be something like

> (iii') In the world where I am faced with a situation very much like[3] the one which faced Winston Churchill, I do not escape from Pretoria.

Such an approach does not really solve the problem, however, because it simply obscures the relation between what Winston Churchill did and who he was. Consider: part of the reason that Churchill was in South Africa in the first place was just that he was politically ambitious, and believed that a high wartime profile would help his political career. Moreover, he used his family connections to ensure that he could get a correspondent's position. Since I have neither political ambition nor powerful family connections, it is obvious that I

---

[3]How much will be "alike enough"? That appears to be part of the difficulty, and will be a problem no matter which analysis we use. Some analyses will be more natural than others, however; see chapter 5.

would not be in a position to escape from Pretoria during the Boer War; but *that* does not capture the sense of (iii) at all. Still, Lewis can say that, in the possible world which makes (iii) true, I have a counterpart who is alive during the Boer War, has powerful connections, has political ambitions, &c. Lewis needs only enough similarity to Churchill that the comparison be reasonable, while yet providing enough extrinsic similarity, and especially, a match of origins, such that the person who does not escape from Pretoria is my counterpart.

Now the trouble is that the counterpart is not at all like *me*. We are, it seems, back where we started. We want the person who does not escape from Pretoria to have enough similarity to Winston Churchill to somehow "count" as comparable to him, while yet having enough similarity to me that we can say it is me. The problem is, the biographies of Winston Churchill and Andrew Sullivan are not adequately similar to allow that. This holds no matter where (or if) we decide to draw the line between intrinsic and extrinsic similarities. In particular, for me to have the connections that Churchill used in order that he be in South Africa, I would have to come from an entirely different background. My parents, then, would also have to be very different from my actual parents. We wind up interpreting (iii) as entailing something like, "If my life had been completely different than it actually is, I would fail to escape from the Boers." That may be true; but it is silly. Certainly, it is nothing like what we take (iii) to be about in English. Whatever (iii) says, it does not seem to have

anything to do with the station of my parents in Victorian England, nor with the date of my birth. It seems to be saying something about my character as compared to that of Winston Churchill -- the actual one. Example (iii), then, is still not easily handled under modal realism. That is not a good reason to reject the example, however; neither would it be sufficient reason to reject modal realism. Nevertheless, it highlights the strange ways in which modal realism must interpret modal sentences in some perfectly normal conversational contexts. That, at least, ought to give us some pause.

3.4 Worlds, empiricism and logic

When we introduced the notion of a world, it was to analyse the notion of something being necessarily true. We found we could understand the notion of necessity by supposing that anything which is necessarily true is true at every possible world. Lewis then suggested that the best way to understand modality is to believe that the possible worlds are real things, which make modal claims true. But it is not clear that we can say anything about the worlds without falling into circularity.

Let us suppose that Lewis is right. Suppose that there are many worlds; and that the truth of 'ϕ' in some of them is what makes 'possibly-ϕ' true in this one. The obvious question is, how can we know whether 'ϕ' is true at any world? Lewis claims that we can know

it (at least in some cases) *a priori*[4]. But Lewis must make a non-trivial assumption if his epistemic defence is to work. He says, for example,

> If modal knowledge is what I say it is, and if we have the modal knowledge that we think we do, then we have abundant knowledge of the existence of concrete individuals not causally related to us in any way. For instance, we know *a priori* that besides the donkeys among our worldmates there are countless other donkeys, spread over countless worlds. (*PW*, p 110)

Lewis offers this in response to the claim that we are not causally related with our knowledge of the other worlds. Normally, I would agree with him; I am willing to suppose that not all knowledge requires causal acquaintance. But in this case, we have a problem; *a priori* knowledge, if it is possible, requires that the laws of logic apply across the worlds. But we cannot have reason to believe that, if we are not causally acquainted with the other worlds. For without *any* such acquaintance, we cannot know that *anything* is necessarily true.

In order for something to be a law of logic, it surely must be necessarily true. It is possible to build a logic from some very minimal laws. But it is not possible to build a logic with no assumptions. In order for us to have a true logic, then, we will need to ensure that the original assumptions are true in all possible worlds. But the worlds are isolated: we "can't get there from here." We can nowise find out whether any given logical assumption is necessarily

---

[4]See *PW*, § 2.4, and below.

true. So, we have no necessarily true rules of logic. That looks like bad news for any *a priori* modal knowledge we might have. Alternatively, Lewis can claim that we know some law of logic *a priori*. But we would be right to ask why we ought to believe him when he claims as much.

Such a view might seem strange; but, consider again what Lewis is saying in his example above. Surely he cannot mean that we have *a priori* knowledge of possible donkeys before we even know of actual donkeys; for that would mean that we must be able to know that something is possible before we know that anything is actual. At the very least, such a thesis would be unverifiable: by the time we can discuss possibilities, we already are aware of too many actual things. It seems likely, then, that Lewis just means that, in knowing of actual donkeys, we also, without experience, know the possible donkeys. They are, of course, unactualised; but, in seeing a given donkey, we can imagine another which is right next to it -- perhaps, even, one that is talking. Now, if we have such *a priori* knowledge, it must be because of our ability to proceed logically from "this donkey" to "possible donkeys"; we can have no experience of the other-worldly donkeys, so our only route to them is from this world, via reason. But if that is the case, the same must be true of the laws of logic, which ought to govern our reason. And there is surely no certainty in assuming our logic to be necessarily true in order to prove that our logic is necessarily true.

One might be tempted to claim that at least some truth of logic is just *obviously* true: say, the Law of Contradiction. I can only

respond: why is that obvious? Once we admit that the truth of modal sentences is dependent upon the states of affairs in other, completely inaccessible realms, what basis do we have for the claim that *anything* is obviously true? Admitting the notions of plenitude and closeness, we can see that some modal sentences will be obviously true: some worlds will be very like this one, and (since they must exist, to fulfill plenitude) that makes certain (truthlike) claims true. But this is a very poor reason to suppose that all of the worlds are like this one, and so gives us nothing like the *a priori* modal knowledge that Lewis says we have.

Perhaps Lewis would be happier to say that some *a priori* knowledge is *strictly a priori*: it is "hard wired" into us. The difficulty now is twofold. First, unless there is some reason why this knowledge could be guaranteed as true, we have no reason to believe in it. But it cannot be guaranteed to be true, insofar as we cannot have a proof that it is true; for that would just be the bootstrapping operation we considered above. Second -- and I take this to be the more serious objection -- such "hard wiring" is really just a device to sweep the analysis problems under the carpet. The original motivation for modal realism was to provide tools whereby we reduce the number of problematic notions that we take as primitive. But, surely, *a priori* knowledge is at least as troublesome as the notion of possibility. *A priori* knowledge which we obtain by dint of birth seems particularly mysterious. It seems, then, that the price of modal realism is higher

than advertised: not only do we get an ontology that (Lewis admits; cf. *PW*, pp 134-135) defies common sense, but we also have to accept a completely mysterious form of *a priori* knowledge to establish the ontology.

Maybe Lewis can argue that I am not paying attention to my experience. He might point out the tremendous success of logic, and its utility in reasoning. It seems not to fail. He can argue, then, that I am not being true to myself: that I am willing to let what I believe in the rooms of philosophy departments conflict with what I believe in daily life. As he argues in *PW*, §2.5, modal realism gives no more reason to suppose that an evil demon is fooling me than everyday induction gives. Scepticism cannot be ruled out by experience before one is a modal realist, but most of us are not (Cartesian) sceptics anyway; so why should we be so sceptical when we accept modal realism?

In arguing this way, Lewis seems to miss that, prior to believing in modal realism, I had no motivation for suggesting that the world really was a place where an evil demon fooled me all the time. Under modal realism, however, I am supposed to believe not only that I *could* be fooled, but that, in some worlds, somebody *really is* so

fooled.[5] So, he says, "Why should the reason everyone has to distrust induction seem more formidable when the risk of error is understood my way: as the existence of other worlds wherein our counterparts are deceived?" (*PW*, p 117) Now, under ordinary scepticism, the question is whether "I am being fooled" is something I should believe. Under Lewis's modal realism, I should already accept that *someone* should believe that statement; the only question is whether that someone is me. Moreover, under modal realism, scepticism is not merely a problem for induction (which is what *PW* §2.5 is really about); rather, we have a live question as to whether we should be Cartesian sceptics. We have no foundation for logic, as we saw above.

In the end, we might want to say that the truths of logic are necessary for our thought; so they must be true. But that is no help. For we wanted the worlds to include every possible way a world might be. We cannot then restrict ourselves to whatever we can imagine a world might be like; our failures of imagination cannot restrict the use of the word "possible", unless we want those failures to be a stricture at all times. And, of course, we would then have two problems: first,

---

[5]The sceptical argument against modal realism was originally proposed (independently) by Forrest and Schlesinger. Schlesinger's objection does not seem concerned with Cartesian scepticism. Forrest, on the other hand, initially considers the case of the evil demon, but then abandons that objection in favour of what he takes to be a more rigourous kind of scepticism. It seems, however, that it is the case of the Cartesian sceptic which is troublesome for Lewis's view.

whose imagination do we use as the standard; and second, why is the "incredulous stare" response to Lewis not a good one?

At issue is something like the problem Descartes ran into when claiming that God created the eternal truths:

> It will be said that if God had established these truths he could change them as a king changes his laws. To this the answer is: Yes he can, if his will can change. 'But I understand them to be eternal and unchangeable.' -- I make the same judgement about God. 'But his will is free.' -- Yes, but his power is beyond our grasp. In general we can assert that God can do everything that is within our grasp but not that he cannot do what is beyond our grasp. It would be rash to think that our imagination reaches as far as his power. (Letter to Mersenne, 15 April 1630: CSM III, p 23).

It seems utterly strange to state that the rules of logic, or the eternal truths, or anything of the kind could be anything other than necessarily true. But if we are to believe that the necessity of a given truth rests upon its being true in every possible world, and that those worlds are completely "cut off" from our inspection, then we have no basis for claiming that anything is necessarily true. We can, of course, assume that something is necessarily true; but, then we will not be able to refute those who claim that we have not included everything in the set of all possible worlds.

Perhaps this merely means that Lewis has to admit of impossible worlds. His original argument, to the effect that "at world w" is a restricting modifier, just like "on the mountain", meant that nothing of the form "possibly, ϕ and not-ϕ" was permissible (cf. *PW*, p 7 n). But if we are agnostic about whether the Law of Contradiction holds at every world, the impossible worlds come to be as respectable as

any other world. But that suggests that modal realism is not the bargain that Lewis offered. In fact, it suggests that the cost of the philosophers' paradise is just too high.

Chapter 4

Existential Quantification and Existence

Lewis claims that the best way to understand talk of *possibilia* is to accept such talk as the literal truth (*PW*, p 4). That claim, at its base, amounts to a claim that whenever an existentially quantified statement is true, it is so just because of its correspondence with some state of affairs. That is a view which, in another form, can be found in the writings of Russell and Quine. Of course, both Quine and Russell are ontologically conservative; nevertheless, Lewis's argument depends on the view that our existential quantifications force us to particular ontological commitments. If that view turns out to be questionable, then we have yet another reason to wonder about the value of modal realism.

4.1. Existential quantification 1: conservatism

It is easy to see why we should suppose that language is, in essence, a 'naming device'. Consider the tempting description offered by Augustine:

> When they named any thing, and as they spoke turned towards it, I saw and remembered that they called what they would point out by the name they uttered... And thus by constantly hearing words, as they occurred in various sentences, I collected gradually for what they stood; and having broken in my mouth to these signs, I thereby gave utterance to my will. (Augustine, *Confessions*, I.8)

This is hardly an unusual description of the way language works. Consider what Quine has to say about the same phenomenon:

> I hold...that the behaviorist approach is mandatory. In psychology one may or may not be a behaviorist , but in linguistics one has no choice. Each of us learns his language by observing other people's verbal behavior and having his own faltering verbal behaviour observed and reinforced or corrected by others. We depend strictly on overt behaviour in observable situations. As long as our command of our language fits all external checkpoints, where our utterance or our reaction to someone's utterance can be appraised in the light of some shared situation, so long all is well. (*Pursuit of Truth*, henceforth *PT*, §14)

Quine's behaviourism finds expression, famously, in a thought experiment in which an English-speaking linguist is to interpret an utterance by a speaker of a natural language, previously unencountered, called "Jungle". The story goes that the Jungle speaker points to something the linguist calls a rabbit, and utters, *"Gavagai!"* The linguist, supposing that "*gavagai*" means "rabbit", tries out "*gavagai*" on occasions where he thinks he sees rabbits, and looks to see if he gets the expected reaction from Jungle speakers. If so, he supposes that "*gavagai*" means "rabbit" (*PT* §16).

Quine takes these sorts of "observation sentences" to be the fundamental basis of any kind of talk about anything.[1] For, he says,

---

[1]Or, at least, anything *real*, whatever that might be. It seems, for example, that Quine has little use for traditional metaphysics at least partially because it is not obviously rooted in some (possibly infinite) set of observation sentences.

> Theory consists of sentences, or is couched in them; and logic connects sentences to sentences. What we need, then, as initial links in those connecting chains, are some sentences that are directly and firmly associated with our stimulations. Each should be associated affirmatively with some range of one's stimulations and negatively with some range....
>
> I call them *observation sentences*.... Unlike 'Men are mortal', they are *occasion* sentences: true on some occasions, false on others....Briefly stated, then, an observation sentence is an occasion sentence on which speakers of the language can agree outright on witnessing the occasion.... (*PT*, §2)

So, we see that observation sentences are the fundamental carriers of meaning. A community of speakers can be (in principle) "infiltrated"; that is, it can be interpreted by outsiders because of the ability to learn to assent to external stimuli in the same way. To explicate further, we might say that the linguist can learn to speak Jungle not because languages are translatable -- for translation is at least underdetermined, if not indeterminate -- but because reference to objects is common to all languages, and what each speaker is trying to do is refer to the world. Observation sentences are crucial, then, because they are what ties language to the world, at least in practice.

The key to this notion of how language refers to the world, however, lies in a view about the status of the objects. Quine claims that these objects are reifications. We get the pure sensory experience, develop a concept of an object "out there", and then posit that the object *is* "out there", causing our sensations. So, the language is really referring to the sensory stimulations: we reify the objects based upon pure sense data. Quine wants to dispense with the reified objects, because of the same sort of underdeterminacy in

ontology as there is in language. Whether "*gavagai*" refers to an object (rabbit), a set of undetached rabbit parts, everything in the universe *except* the rabbit, or the condition of the universe such that "it's rabbiting" (along the lines of "it's raining") makes no difference on this view, because the logical structure remains the same no matter what sorts of reification we come up with. That is, both utterances, "Lo! a rabbit!" and, "It's rabbitting!" find logical expression in the same way: $(\exists x)\ (Rx\ \&\ (\forall y)(Ry \supset y=x))$.[2] The reification that a given language makes is useful; but the real ontological commitments are found in the underlying logic. So, even though a particular speech-act -- whether it be an observation sentence or some other sentence -- may reify a number of objects, the speaker's "real" ontology flows not from that reification, but can be determined by figuring out the underlying logic.

We can see, then, that Quine's view is really a refinement of the theory of language offered by Augustine. The difference is really a matter of how "hard" we imagine the ontology of a given speech-act to be. Augustine's view is, as Wittgenstein says, that

> the individual words in language name objects -- sentences are combinations of such names.-- In this picture of language we find the roots of the following idea: Every word has a meaning. This meaning is correlated with the word. It is the object for which the word stands. (*Philosophical Investigations*, henceforth, *Investigations*, §1)

---

[2]It might be argued that "here and now" need to be accounted for. I have left out space-time indexicals for the sake of simplicity.

Quine's view moves the meaning from words to observation sentences: we only analyse observation sentences word by word "from the retrospective standpoint of theory." (*PT*, §10) The real difference here, however, is not great. For both Augustine and Quine take some (comparatively small) set of building-blocks, and try to make the rest of the language out of those blocks. Quine simply has a more sophisticated mechanism for interpreting the underlying assumptions. Moreover, in both cases one is referring to the world; the difference is in whether the objects are "real things" (as in Augustine's case), or whether they are reified from sensory stimulations.

That Quine should want to avoid positing a direct connection between words and objects is hardly surprising, for in his notion of ontological commitment we can see the influence of Russell's theory of descriptions. Recall that Augustine offers an account of how words are related to objects. I have argued that Quine refines such a theory. These refinements are rooted in the concerns which originally led Russell to concern himself with language. For Russell's problem was really to show that psychologism was a poor foundation for mathematics. When he discovered Peano's symbolic logic, he thought to deduce the whole of pure mathematics from logical principles. In order to do so, he had to define what he called "the indefinables of mathematics": to explain "the fundamental concepts which mathematics accepts as indefinable" (*Principles of Mathematics*, henceforth *PoM*, xv). Significantly, these indefinables include naming and denoting.

Russell's account, as offered in *PoM*, leads to some intractable difficulties because of the device of denoting concepts. Russell's solution to this was to abandon denoting concepts in favour of quantified variables and predicates. In most cases, he thought, the matter was not at all difficult: the cases of 'a', 'some', 'every', 'all' and 'no'. The challenge arises in the case of 'the'.[3]

Russell considers an instance: "The father of Charles II was executed." (p 417) He claims that 'the', strictly used, is an expression of uniqueness: "[W]hen we say '*x* was *the* father of Charles II' we not only assert that *x* had a certain relation to Charles II, but also that nothing else had this relation." ("OD" p 417) So, our original statement becomes

> "It is not always false of *x* that *x* begat Charles II and that *x* was executed and that 'if *y* begat Charles II, *y* is identical with *x*' is always true of *y*". ("OD" p 417)

This analysis turns out to be an entailment of the principle of the theory of denoting which Russell offered:

> that denoting phrases never have any meaning in themselves, but that every proposition in whose verbal expression they occur has a meaning. The difficulties concerning denoting are, I believe, all the result of a wrong analysis of propositions whose verbal expressions contain denoting phrases. ("OD" p 416)

---

[3]I have restricted myself to the discussion in "OD" for the sake of simplicity; but all of what appears there finds formalisation in *Principia Mathematica*.

We can see that the full expression of "the father of Charles II was executed" dispenses with the direct denotation of the individual in question ('the father of Charles II'), and replaces that with an analysed description of that individual in terms of variables and predicates ('*x*', '*y*', 'begat'). The result, then, is a generalised way of interpreting statements about the father of Charles II, such that any statement "*C*(the father of Charles II)", implies

> It is not always false of *x* that 'if *y* begat Charles II, *y* is identical with *x*' is always true of *y*,

which is usually stated by saying that Charles II had one and only one father. So, of course, if the condition fails, any statement of the form "*C*(the father of Charles II)" is false. That is the reason why, famously, "The present King of France is bald," is false. For "if we enumerated the things that are bald, and then the things that are not bald, we should not find the present King of France in either list." (p 420)

Russell wants to argue that propositions do not become nonsense just because they contain denoting phrases with no denotation. Rather, he argues, a statement like, "The present King of France is bald," is "not nonsense, since it is plainly false." ("OD" p 419) So, he goes on to say, "we must abandon the view that the denotation is what is concerned in propositions which contain denoting phrases." The way to such abandonment, he claims, is the device of primary and secondary occurrence (this is usually known as the scope distinction). We can therefore explain why 'The present King of France is bald' is false:

> If "*C*" is a denoting phrase, say "the term having the property *F*', then "*C* has the property ϕ" means "one and only one term has the property *F*, and that one has the property ϕ". If now the property *F* belongs to no terms, or to several, it follows that "*C* has the property ϕ" is false for *all* values of ϕ. ("OD" p 424)

That means that "the present King of France is bald" is certainly false, as there is no term (i.e. there is nothing) which has the property of being-the-King-of-France-at-present. Moreover, "the present King of France is not bald" is also false if it means

> "There is an entity which is now King of France and is not bald",

but true if it means

> "It is false that there is an entity which is now King of France and is bald." ("OD" p 425)

In other words, "the present King of France is not bald" is true just in the case where "the present King of France" is a denoting phrase with secondary occurrence.

Obviously, Russell's strategy in interpreting statements about the present King of France is explicitly empiricist. How do we evaluate statements which include a reference to some particular object? We evaluate them by checking them against the world. If we want to know whether the present King of France is bald, we first must go out and enumerate "the things that are bald, and then the things that are not bald..." ("OD", p 420) The truth of a statement is determined by some sort of correspondence to the way the world is. But that seems to lead us to grief when we consider objects which everyone knows to be non-existent. For it seems we need to find out whether any given denoting

phrase has a denotation. When there is no denotation, a proposition containing the denoting phrase will be false (in the same way 'the present King of France is bald' is false). Russell seems to have thought, as Cassin argues (pp 278-80), that the matter was one of empirical investigation. If we want to know whether there is a present King of France, we go out and look for one, or (in this case) find out about the Constitution of France in order to discover whether there might be a King of France.[4] There are objects, however, about which we can speak coherently, and about which we can say true things, even though those objects do not exist. Most of these are *possible* objects: their existence would not be self-contradictory. Examples include the unicorn, Dr Juvenal Urbino[5], and my first-born son. For each of the following is true:

> The unicorn has one and only one horn.
>
> Dr Juvenal Urbino saw his friend, Jeremiah de Saint-Amour, dead of cyanide poisoning.
>
> My first born son will always be related to me.

---

[4]It is interesting to note, of course, that this makes our logic badly underdetermined, since we have no idea whether a given denoting phrase should *really* be treated as a secondary occurrence. That is, even though the logical structure is well-developed, we will have an intractable epistemological problem built deeply in our theory of denotation. This is also the reason why "quasi-real" objects like unicorns are going to be hard to account for under Russell's view.

[5]A character in Gabriel García Márquez's *Love in the Time of Cholera.* The usual fictional character for such an example is Sherlock Holmes; but, poor Sherlock has of recent years come to be something more like the unicorn, appearing in guises which vary slightly from author to author. See note, below.

Russell's view deals easily with the last of these: we need only time-indexation or an analysis of the necessary truth of the claim, or both. The first two, however, seem more contentious. One of them involves a denoting phrase which does not denote anything real in the world, but about which there is much agreement as to its properties; the other involves a denoting phrase which is true just in case we specify the domain in which it might be true (i.e. the story in which the character appears). In other words, everybody who knows what a unicorn is knows both that unicorns have one and only one horn *and also* that unicorns do not exist in any forest, enchanted or otherwise, in the world. Similarly, anyone who has read *Love in the Time of Cholera* will recognise the truth of a description of the opening scene of the book. Nevertheless, neither of the first two statements, above, can be construed as true on Russell's account of denotation.

We could try many approaches to resolving the problem of what we might call "conventional objects", like unicorns, and "fictional objects", such as Dr Urbino. It seems to me to be foolish to deny the existence of these things, since we seem able to do something like empirical investigations about their natures. If we want to know about unicorns, we find that there is a well-developed literature about them; part of that literature notes that they are not things that actually exist. Similarly, if we want to know about Dr Juvenal Urbino, we can go to García Márquez's book, and look there for the description of the eminent doctor. Of course, there are things we cannot find out about

these objects: the internal anatomical structure of the unicorn, say, or the second-favorite toy of Dr Urbino when he was seven and a half years old.[6] Nevertheless, we can, it seems, properly assert at least that the unicorn has one and only one horn, and that Dr. Urbino faced the death of his friend. Now, there seem to be two obvious strategies to deal with these cases under Russell's view. The first is to try to specify a domain such that the objects of the denoting phrases exist, possibly in some "honorary" way[7]. So, some claims about Dr Urbino are true-in-a-story, and some claims about the unicorn, conventionally so. Such an approach has the awkward result of dividing the world up into rather *ad hoc* classes of kinds of existence; and, though we may have gradations in our ontology, it is not obvious that these gradations arise because of our semantics: the utterance, "Dr Urbino was a well-respected man" is true in English regardless of whether anybody has read *Love in the Time*

---

[6]It is also important to note that some objects tend to straddle the rather artificial line I have drawn: Sherlock Holmes is, in one sense, a fictional character in books by A.C. Doyle; at the same time, Holmes conventionally wears a double-peaked cap, because every time he appeared on television or in movies, he wore such a cap.

[7]See. e.g., Woods, who suggests that the real solution to any of these problems is just to provide the right sort of semantics. In particular, he illustrates the differences between "non-entities" and "nonesuches", then considers the claims we might make in favour of truth-by-convention. Woods' approach is clear, well-written, and certainly enlightening; but, it nevertheless depends upon language being essentially a naming device, which also allows for the parasitic use of language for metaphors, etc. The trouble seems to me just that some kinds of existence will get a strange sort of "priority" over others. On this view, everything is real; but, some things are more real than others.

of Cholera[8]. Since all suggestions for making propositions about fictional objects true which rely upon truth-by-convention depend upon somebody knowing the convention, or upon a "sayso" semantics which relies upon the author's intentions (Woods's proposal), we seem stuck with truth that is nothing more than whatever somebody says. In other words, under this approach, I ought to be able to change the make-up of unicorns simply by writing a new story. That seems wrong: we would know something to be awry were we to encounter a story about a two-horned unicorn.

Woods has also offered reasons to suppose that alternative approaches to truth in fiction fail; so, the second strategy is to suppose that there are realms of existence just like ours *except that* the relevant extra bit (say, unicorns) gets added to (or, in some cases, taken from) existence. This latter is, of course, the strategy of modal realism.

Modal realism, then, is an extension of a fairly common empiricist strategy in interpreting language. If there are problems with that empiricism, however, then those problems will extend to modal realism. At the very least, if there is another view which we find plausible, we may conclude that modal realism is not our only option.

---

[8]Although probably not "written"; for if there were no standard of truth for claims about Dr Urbino (as, for example, the book provides), it does not seem we could plausibly say anything true (or false!) about him.

4.2 Correspondence and truth

> Seen in retrospect, the failure of correspondence theories of truth based on the notion of fact traces back to a common source: the desire to include in the entity to which a true sentence corresponds not only the objects the sentence is 'about' (another idea full of trouble) but also whatever it is the sentence says about them. (Davidson, "True to the Facts", p 49)

We saw, above, that views like those of Quine and Russell require that the test of the truth of some sentence be whether that sentence is confirmed by the actual state of affairs. That provision amounts to a requirement that a true sentence be "true to the facts." But what are these facts? The more we consider what they might be, the more obscure they become.

The most obvious difficulty is that it is next to impossible to describe a fact except with the sentence that the fact is supposed to verify. As Davidson notes, "...the fact that verifies 'Dolores loves Dagmar' should somehow include the loving. This 'somehow' has always been the nemesis of theories of truth based on facts." ("True to the Facts", p 48) Moreover, the notion of fitting the facts is not really helpful. It is no better than simply saying that something is true.

To see this another way, consider Russell's example of the present King of France. He claims that we would not find the King of France among the list of all the bald things; neither would the King appear in the list of all the non-bald things. It seems, though, that there is nothing more to that claim than that there is no King of France. For we have no more knowledge of the list of all the bald

things, and all the non-bald things, than we have of the non-existence of the present King of France.

The notion that sentences are made true by facts is rooted in the distinction between scheme and content. As Davidson observes,

> The idea is then that something is a language, and associated with a conceptual scheme, whether we can translate it or not, if it stands in a certain relation (predicting, organizing, facing, or fitting) [to] experience (nature, reality, sensory promptings). The problem is to say what the relation is, and to be clearer about the entities related. ("On the Very Idea of a Conceptual Scheme", p 191)

The scheme is supposed to fit the experience in one of two ways. Conceptual schemes, or languages, might "organise" experience, or they must "fit" the experience. (In the case of the views under consideration, the conceptual scheme should "fit" experience.) Indeed, the thing to which the scheme must fit may be of two kinds: "reality" (e.g., in Russell, the language must fit the reality), or "experience" (e.g. in Quine, the language of observation sentences arises from the sensory promptings). But in any case, the notion is that the language must somehow "map on" to something external to it: reality or experience. But

> [t]he trouble is that the notion of fitting the totality of experience, like the notion of fitting the facts, or of being true to the facts, adds nothing intelligible to the simple concept of being true. To speak of sensory experience rather than the evidence, or just the facts, expresses a view about the source or nature of evidence, but it does not add a new entity to the universe against which to test conceptual schemes. ("On the Very Idea of a Conceptual Scheme", pp 193-4)

What alternative might there be to this notion of a scheme which fits experience? What we want, of course, is a way to understand utterances. Davidson suggests that we adopt something like a theory of truth which is relativised to occasions of speech, combined with a strong notion of translation. So, a theory of truth would define a three-part predicate, "T *s,u,t*", where *s* is a sentence, *u* is a speaker, and *t* is a time. We say, "'sentence *s* is true (as English) for speaker *u* at time *t*'..." ("True to the Facts", p 44) Of course, this account leaves us with the difficulty of providing a detailed semantics of natural languages, and developing an account of translation which is not dependent upon the concept of meaning. It leaves us with a different set of problems in saying what truth consists in. But such a set of problems is not obviously more intractable than the problems we face in understanding how a sentence might "fit the facts". I do not know which set of problems to prefer in the case of interpreting sentences about the actual world. But Lewis uses the same sort of scheme-content distinction in the service of specifying the truth-conditions of modal sentences. We should explore whether the difficulties around scheme and content are more troublesome when we begin quantifying across worlds.

4.3 Existential quantification 2: liberalism

We wanted to see whether the scheme-content distinction will cause trouble for modal realism. First, it will be useful to reconsider the outline I have offered, and then note a deep assumption -- one which, I think, is untenable. We noted, to begin with, that the notion

that language is essentially a naming device is, at least, plausible. A simple description of how people come to learn a language seems to be enough to show us that language names things. We also observed, however, that Quine (for reasons which we did not pursue) developed a sort of "naming" theory of language which did not depend upon a 1:1 correlation between a name and its named object. Instead, he offered a view which depended primarily upon observation sentences. These observation sentences can be translated into quantified logical notation in order to dispense with certain kinds of ontological difficulties stemming from the reification of the "objects" named by the observation sentences. We saw that Quine's approach amounts to pushing the ontological troubles to a deeper level, without neutralising them; for this account eliminates the reified objects by making their existence a matter of whether they qualify as the value of a bound variable. That approach is, at its root, the theory of Russell's "OD". And, we discovered, Russell's approach makes false some propositions that everyone takes to be true. I then suggested that one way to circumvent this difficulty is to understand the purported truth of such propositions as amounting to their *possible* truth, and *actual* falsity. It is here, apparently, that modal realism gains plausibility. But, it is only plausible because of a deep confusion about the nature of language. That confusion is expressed nicely by Russell, in *PoM*, in a passage which he takes to be expressing an obvious truth:

> [I]t must be admitted, I think, that every word occurring in a sentence must have *some* meaning: a perfectly meaningless sound could not be employed in the more or less fixed way in which language employs words. The correctness of our philosophical analysis of a proposition may therefore be usefully checked by the exercise of assigning the meaning of each word in the sentence expressing the proposition. (p 42)

In "OD", of course, Russell supposes that he has abandoned the above to some extent. Certainly, Quine does not suppose that he holds such a difficult view. But as I argued in section 4.2, both authors imagine some way of mapping "language" -- even indirectly -- onto "the world". In such a view, there is some "tight" relation between the utterances of a speaker (if they are not nonsense) and the world about which the speaker attempts to express some proposition. The relation works two ways: "the world" causes stimulations in the speaker, which result in the utterance of a proposition. The truth of the proposition is then tested by the world. For Russell, the latter, "testing" part is more direct than for Quine. In Quine's case, the decision about the truth of a proposition is to be tested by other speakers under relevantly similar conditions. But, regardless of whether we make our ontology relative (and indeterminately translatable, as does Quine) or externally determinate (as does Russell), we find that the external stimulations are supposed to be the cause of more-or-less fixedly interpretable speech acts of individuals who share a language. That is, there is a common "external world", which the speakers share, and which works as a foundation for meaning in language. All that modal realism adds to such a mix is the realm of *possibilia*. That is to say, Lewis's contention that "systematic philosophy goes more easily if we may presuppose modal

realism in our analyses" (*PW*, p vii) really means that, if we want to understand utterances which include our modal idioms, we have to suppose that those modal idioms are naming something, somewhere. They are existential quantifications over worlds. I want to suggest that these existential quantifications ought not to make our ontology: reference to "real objects" is not *the* criterion of truth.

Recall that Lewis's argument for the existence of possible worlds is that any claim that things "could have been otherwise" amounts to an existentially-quantified statement. That means that there is something which is the value of a variable, and that thing is a possible state of affairs or, as Lewis prefers to call it, a possible world. But this means that Lewis includes in his theory Davidson's "third dogma of empiricism": the scheme-content distinction. What is peculiar in Lewis's case is that included in the "content" are things which are deemed to be "reality" even though they are not -- indeed, could not be -- part of anyone's experience.[9] Regardless, we can see how the distinction functions. There is a theoretical scheme, which for Lewis is language including modal operators. The language is thought to stand in a certain relation -- whether one of "fitting" or "organising" -- to

[9]At least, anyone in this world. Of course, for Lewis, the inhabitants of the alternate world have all the experiences of that world. We have already seen (chapter 3) some difficulties arising with some individuated counterparts -- e.g., the case of me-as-Churchill. Also, as we saw in chapter 2, it is not enough that these alternate states of affairs are "objects of thought", unless those thoughts correspond to some external other-worldly reality; otherwise, we would have some brand of ersatzism.

the "reality". The reality, however, includes every non-self-contradictory object and state of affairs to which the language can refer (this, due to logical plenitude). Lewis might want to claim, then, that he does not face the normal sort of "third dogma" objection, just because the content and the scheme are not separable. Such a defence fails, of course, just because the original justification for treating possible worlds as real was that such treatment makes sentences of the form 'possibly-ϕ' true whenever 'ϕ' is logically possible. In other words, the possible worlds are one more part of "reality", and it is correspondence to that reality which makes sentences in the language -- the scheme -- true.

One might like to suggest, on behalf of Lewis, that the above is an unfair treatment, since Lewis never actually claims outright that 'possibly-ϕ' is true just in case the statement 'ϕ' corresponds to some fact in some world. But what does it mean to say , "'possibly-ϕ' is true just in case 'ϕ' is true in some world," except to say that 'ϕ' corresponds with the state of affairs in some world? Irrespective of how relativised to a speaker the truth conditions of 'ϕ' are, the truth conditions are either tied to the set of all worlds, or they are not. If they are not tied to any world, the motivation for modal realism disappears; so, Lewis must be dependent upon some kind of "content" to make any statement true.

To see why that is true, consider the case where the content of a world does *not* determine the truth of any statement about it (we

can pick any theory of truth we like which is not a correspondence theory of truth). The speaker's claim is then either true or false not according to the truth or falsity of the statement in any given world, but according to whatever criteria are specified by the alternative theory of truth. For example, if we want to argue in favour of some sort of "coherence" theory of truth, then the truth of 'possibly-ϕ' is a question of whether that statement coheres with other beliefs, propositions, speakers' beliefs, etc. *irrespective of the speaker's world*, and not whether there are states of affairs of some world which make 'ϕ' true in that world. Of course, it makes no difference what makes 'ϕ' true in some world: we might be coherence-theory believers within worlds. Correspondence must, nevertheless, hold *across* worlds if the truth of 'ϕ' in some world is to make 'possibly-ϕ' true in another. But, of course, that means that 'possibly-ϕ' is true if and only if it "fits" the totality, not of experience, exactly, but of reality.

Now, as we suggested in section 4.2, "the notion of fitting the totality of experience...adds nothing intelligible to the simple concept of being true." ("On the Very Idea of a Conceptual Scheme", pp 193-4) Similarly, the correspondence of "possibly-ϕ" to the reality of ϕ in some possible world tells us nothing more than that ϕ is possible. But, at least in the case of Lewis' possible worlds, there is no question as to the possibility of ϕ anyway, since Lewis allows all and only logically possible states of affairs to constitute the possible worlds (see chapter 2). In other words, to say that "ϕ is possible" is

true just in case "ϕ" is true at some world W is, by definition, to say nothing more than "ϕ is possible". Since Lewis takes the reality of possible worlds to be foundational to his analysis of counterfactual statements, there is an obvious difficulty.

The real trouble for Lewis is the same one that gave us pause with Russell and Quine. Recall that Russell runs into problems with 'the unicorn' because, when we enumerate all the one-horned things in the world, and all the non-one-horned things, we do not find the unicorn among either group. What would make the proposition, 'The unicorn has one and only one horn,' true, would be an actual one-horned unicorn. We can consider a comparable possible-worlds example: 'If there were unicorns, they would have one and only one horn.' Such a proposition is true if and only if there is a world such that there are unicorns (who have one and only one horn). That is the case because, in using the device of possible worlds, we simply widen the domain over which we may quantify. In other words, we include in the realm where we do our enumeration not only the actual world, but every other world as well. So, the real things are the things which are the value of a bound variable, although in this case the variable ranges across worlds.

Lewis might claim, by way of response, that his view offers us resources to explain the contextual sensitivity of a natural language. As we saw in section 2.5, Lewis wants to analyse language by relativising it to "linguistic communities", much as interpretation under Quine's view is to be handled by the relevant class of speakers

under relevantly similar circumstances (see *PW*, pp 40 ff, and above). But he suggests we do this by specifying a finite vocabulary, and specifying (for each element of the vocabulary, at every possible world) a syntactic category and a semantic value. But how are we to do that? As Davidson noted, it is not at all clear that we can do so for any given speaker. At best, then, Lewis's analysis forces us to be agnostic about which world a particular speaker is in when we assign the categories and values to his vocabulary. In what way is that better than simply saying that we need to relativise our interpretation manual to our situation? It seems the only advantage is that we have some fact, or set of facts, to which our interpretation manual corresponds. In this, Lewis might be able to get around some of the "reference" problems that Russell has; but, as we saw earlier, all Quine's approach really does is to hide the ontological problems.

Obviously, Lewis's approach is immune to the objection that it makes false a perfectly true statement about unicorns. However, the view leads us to posit the existence of many entities which we say *do not* exist. While Russell explains why "The unicorn does not exist" is true, he has to say that "The unicorn has one and only one horn" is false. Conversely, Lewis can say that the latter statement is true; but, he has to say that "The unicorn does not exist" is, strictly speaking, false. All he can do is patch this over by saying that the unicorn does not *actually* exist; but, it exists nevertheless, just as many other non-actual things exist. And, it is here that there is

something unnatural about the ontological consequences of Lewis's realism. For we are likely to say, "Yes, but the unicorn does not exist *really*." And the root of this problem is the same for Lewis as it is for Russell: if one takes "existence" to be an univocal term, one is forced either to posit all manner of entities in many different worlds, or to deny that meaningful things can be said about things which do not actually exist. Surely, though, the problem presents itself only when we fool ourselves into thinking that a word like existence has one and only one meaning.

One might be inclined to argue that existence need not have an univocal meaning for Lewis. But it is the existential quantification (across worlds) which is to be tested in order to find out whether 'possibly-ϕ' is true. Moreover, Lewis provides a reason to suppose that the meaning of existence is not different from world to world:

> I do not have the slightest idea what a difference in manner of existing is supposed to be. Some things exist here on earth, other things exist extraterrestrially, perhaps some exist no place in particular; but that is no difference in manner of existing, merely a manner in location or lack of it between things that exist. Likewise some things exist here at our world, others exist at other worlds; again, I take this to be a difference between things that exist, not a difference in their existing.... If I am right, other-worldly things exist *simpliciter*... And if I am wrong, other-worldly things fail *simpliciter* to exist. They exist, as the Russell set does, only according to a false theory. That is not to exist in some inferior manner -- what exists only according to some false theory just does not exist at all." (*PW*, pp 2-3)

Obviously, Lewis is of the view that there is only one way for something to exist. If that way fails, the thing does not exist at all. Existence, then, is an univocal term.

What difference does it make if existence is *not* an univocal term? To begin with, if we were to accept that there might be different ways for something to exist, we would not need to worry about the reality of non-actual objects. We could use our modal idioms without supposing that they somehow refer to entities in other worlds. We would then have no real reason to be modal realists; and we would not need to be troubled by the ontological worries that we get with Lewis's modal realism. Of course, such an approach may also mean that the analysis of language, which is a strong point in favour of Lewis's view, becomes much more difficult. In the next chapter, I shall suggest why I think Lewis's analysis of language is not the virtue it appears to be. We have already laid the foundation of that suggestion, however; for as we have seen, the use of existential quantification to analyse language raises as many difficulties as it solves. We may know why 'The unicorn has only one horn' is true; but we will be hard pressed to say why 'The unicorn exists' must be true. Our problem is a simple one: the proposition is *not* true, but false. A theory which makes it true must have something wrong with it.

4.4 Do *possibilia* really help?

What alternative can we have to Lewis's modal realism? We need something that will avoid the problems of reference across worlds;

and, of course, we want something which does not fall into the trap of the scheme-content distinction. What we really want is something like what Davidson suggests. As with any other sentence, we want a modal sentence *s* to be true for a speaker *u* at time *t*. In chapter 5, I will try to suggest some ways of interpreting modal talk with just that sort of schema. But before we proceed, we should ask whether talk of *possibilia* is really as helpful as Lewis suggests.

I have argued that Lewis's modal realism depends on an untenable distinction between scheme and content. The trouble is similar to a difficulty in the views of Quine and Russell. In that case, we seem to have a stand-off: we can take the difficulties of tying the truth of an utterance to the world (by sticking with Russell and Quine), or we can prefer the sorts of problems that arise when we follow Davidson's approach to the matter. But the difficulty is more obvious in Lewis's case. For, whereas we might have some clear idea of what the actual world is like, we are in deeper water when we start considering other, possible worlds. As I argued in chapter 3, we cannot really say anything about the other worlds. It might even turn out that, in some worlds, the laws of logic or mathematics are false. Now, the advantage in talk of *possibilia* was supposed to be that we would get a good theory of possibility; and, it would include (most of) the things which "common sense" says are true, without including (too many) things that "common sense" says are false. So, we have worlds, the contents of which our language is supposed to fit. But, if we have no way of learning those

contents, then we have nothing on which to base our assent to some claims (say, "there could be a talking donkey") while yet rejecting others (say, "there could be a square circle"). The purported reality of the worlds (and their contents), then, does not help to clarify our understanding of possibility. In other words (as we saw in section 4.3), saying that something is possibly true because it is true in some possible world is to say nothing more than that it is possible. For the sake of the very theoretical economy that Lewis likes to invoke (see, e.g. *PW*, p 134), we should simply accept the notion of possibility as primitive in our language. But if we do that, we have lost the motivation for Lewis's modal realism.

Chapter 5

Modal Sentences and Analogy

In the last chapters, we saw some reasons to think that modal realism does not offer us a robust understanding of our modal idioms. In chapter 3, I argued that modal realism cannot easily account for the more difficult cases of possibility. In chapter 4, I argued that the problem is "built in" to the analysis of language upon which modal realism depends: Lewis's view depends too heavily on existential quantification, and is therefore susceptible to any criticism which can be levelled at Quine or Russell. A proponent of Lewis's view might nevertheless argue (justifiably) that none of my critiques are completely devastating. Undoubtedly, we could modify Lewis's view in subtle ways in order to forestall my criticism. In this chapter, then, I want to argue from a different direction. Lewis's reasons for accepting modal realism are partly pragmatic: the theory is serviceable. What follows is an attempt to show that other approaches also capture the meaning of modal idioms; so, they are at least as serviceable as Lewis's theory.

5.1 Trouble in paradise

Lewis's arguments for modal realism work in two ways. The bulk of his argument (as we saw in chapter 2) is devoted to showing how modal realism works (and, thereby, *that* it works). The other argument

is a simple one: the usefulness of modal realism is a reason to think it is true

> Why believe in a plurality of worlds? -- Because the hypothesis is serviceable, and that is a reason to think that it is true. The familiar analysis of necessity as truth at all possible worlds was only the beginning. In the last two decades, philosophers have offered a great many more analyses that make reference to possible worlds, or to possible individuals that inhabit possible worlds. I find that record most impressive. I think it is clear that talk of *possibilia* has clarified questions in many parts of the philosophy of logic, of mind, of language, and of science -- not to mention metaphysics itself.
>
> ...
>
> As the realm of sets is for mathematicians, so logical space is a paradise for philosophers. We have only to believe in the vast realm of *possibilia*, and there we will find what we need to advance our endeavours. We find the wherewithal to reduce the diversity of notions we must accept as primitive, and thereby to improve the unity and economy of the theory that is our professional concern -- total theory, the whole of what we take to be true.... Modal realism is fruitful; that gives us good reason to believe that it is true. (*PW*, pp 3-4)

Immediately after the above, however, Lewis offers a proviso: "Good reason; I do not say it is conclusive." (*PW*, p4) We might have doubts about whether the ontological costs of modal realism are worth the benefits. Or, we might wonder whether modal realism is as beneficial as Lewis seems to think it is; these are the strategies I adopted in chapters 3 and 4. As we saw there, we have reason to think that modal realism is not a good tool to analyse all modal idioms. If that is the case, then perhaps we ought not to believe in worlds.

Nevertheless, it is easy to see that the notion of a possible world is useful. A possible-worlds semantics is very helpful in

analysing modal language. Lewis claims that the utility of believing in those worlds gives us good reason to believe that they exist. But what if we had a view which allowed us the free use of the language of possible worlds, without the ontological commitment? This would be "paradise on the cheap", as Lewis calls it (*PW*, chapter 3). And, Lewis argues, it can't be had; ersatzism will not work. It may be, however, that Lewis has not considered another possibility: that his description of ersatzism is not all there is. Perhaps what we need is a naturalistic approach to language. What we need to do is examine the intent of a speaker before analysing his modal sentences. We can couple that to an ersatzist modal theory which is not subject to his objections.

5.2 Modality as natural speech

Recall the examples I introduced in chapter 2:

(i) I could have become a lawyer.

(ii) I might go to the store later.

(iii) Had I been Winston Churchill, I would never have escaped from Pretoria.

We found that (i) and (ii) were easily analysed under modal realism. But we also had some difficulty in understanding (iii). Still, (iii) is plainly not meaningless. There is *something* we are trying to communicate when we say it. The problem with the possible-worlds analysis is that it depends too heavily on the existence of other-worldly entities. When we run into statements like (iii), we cannot

understand what such entities would be, since anything that is enough like Winston Churchill is not very much like me (and vice-versa).

If we want to understand the examples, we should have something to say about interpreting language. The trouble is, ordinary language is not neat. It does not fit easily in a regimented formal language. If we cannot capture all of our modal idioms in a regimented language, then we will have to equip our language "either with outright ambiguities, or else with devious rules which look at what a formula says before they know what it means to satisfy it." (*PW*, p 12) Lewis thinks that we can get around the difficulty by using modal realism to interpret things directly (cf. *PW* pp 6-14). But we do not need modal realism to interpret the language. Modal idioms mean different things under different circumstances; and, a competent user of language knows that perfectly well.

Part of the problem in trying to offer a theory of how we should interpret modal idioms is that they are very context-sensitive (as Lewis argues: see chapter 2 sections 2.2 and 2.3). If I say, "I might go to the store later," without context, it makes no sense. But consider the following exchange:

A: "Oh, I forgot to buy milk!"

B: "Well, I might go to the store later."

Implicit in B's response is a question: something along the lines of, "Would you like me to pick up milk while I'm at the store?" The utterance makes sense in the context. Now, of course, such an

interpretation is consonant with a modal-realist interpretation of B's utterance. But, we can *just as easily* understand the utterance as a combination of two things: B's belief in B's freedom to go to the store, and an offer to help A by getting some milk. There is nothing strange in claiming that "I might go to the store later" is just a statement of belief about one's freedom to do as one wishes in the future. And, that interpretation does not lead us to posit strange realms, in some of which I go to the store, in some of which I do not. We can understand the sentence perfectly well, and we do not need any of the unhappy ontological baggage that the modal realist must carry.

It becomes more obvious that contextual sensitivity is important when we consider other sorts of exchanges which include the same sentence. Consider:

> A: Would you do me a favour? Could you pick up some tomatoes, basil, and vinegar later?
>
> B: Well, I *might* go to the store later...

In this case, B's response is, effectively, an attempt to get out of doing A's errand. While the sentence has not changed, the meaning has: we could as easily interpret the utterance as meaning, "I really cannot guarantee that I'll get to the store to run your errand." We can imagine, for example, that B may be working late, and so may not get to the store before it closes. In any case, B is expressing a disinclination to do A's errand, as much as B is expressing a belief (or doubt) about his own future.

The same example, (ii), could be a sign of insanity or feeble-mindedness under still different circumstances. If we are faced with a bedridden person suffering from dementia, we will take the utterance of, "I might go to the store later," as yet more evidence that the person is confused. Of course, under modal realism, the statement would be, strictly speaking, true: there is a world in which the person's counterpart is perfectly well, and goes to the store. But, that is not *our* world. In our world, the statement is almost nonsense; it makes no more sense than the utterance, "We're going to storm the beach now," when uttered in the security of a suburban living room. Lewis can, of course, deflect this criticism by claiming that, in this case, there is an implicit restriction of the domain: the interpretation ought to be restricted to this world. And, of course, we *can* interpret the language that way; but why should we? For that is not the only way we can get talk of *possibilia*.

The reason a modal realist approach to modal idioms fails is that it takes the idioms to be referring to a state of affairs; that is the nature of existential quantification. The case of example (ii) shows, however, that some instances of modal sentences are really expressing belief. This means that, when a statement like (ii) is true, it is not true because of the truth of "I go to the store" in some world. It is true, rather, because it expresses what William James calls a living option:

> Let us give the name of *hypothesis* to anything that may be proposed to our belief; and just as the electricians speak of live and dead wires, let us speak of any hypothesis as either *live* or *dead*. A live hypothesis is one which appeals as a real possibility to him to whom it is proposed. If I ask you to believe in the Mahdi, the notion makes no electric connection with your nature, -- it refuses to scintillate with any credibility at all. As an hypothesis it is completely dead...
>
> A living option is one in which both hypotheses are live ones. If I say to you: "Be a theosophist or be a Mohammedan," it is probably a dead option, because for you neither option is likely to be alive. ("The Will to Believe", pp 717-18)

The reason the first two cases of (ii), where A offers to pick up some milk, or hints that he'd rather not run B's errand, make sense to us is that they are live options: A believes that he can go to the store, even if the likelihood is very small (as in the second case). But in the third case of (ii), where the invalid supposes that he might go to the store, we recognise that the statement is false because we recognise that such an event is not at all a live option. Given the state of affairs that actually obtains, the person is not going to be able to go to the store. We might claim that this still implies *possibilia*: a "live option" is still a possibility. But it seems that the important matter, in these cases, is belief in the possibility, rather than the possibility itself.

It is not at all surprising that we cannot properly interpret utterances without considering the conditions of utterance. Language use is, after all, a human activity; such activities do not occur in the abstract, but only in relation to one another. So, we cannot understand an utterance without its context.

To see that more clearly, consider what Wittgenstein says:

> 135. But haven't we got a concept of what a proposition is, of what we take "proposition" to mean? -- Yes; just as we also have a concept of what we mean by "game". Asked what a proposition is -- whether it is another person or ourselves that we have to answer -- we shall give examples and these will include what one may call inductively defined series of propositions. *This* is the kind of way in which we have such a concept as 'proposition'. (Compare the concept of a proposition with the concept of a number.)
>
> 136. At bottom, giving "This is how things are" as the general form of proposition is the same as giving the definition: a proposition is whatever can be true or false. For instead of "This is how things are" I could have said "This is true". (Or again "This is false".) But we have
>
> 'p' is true = p
>
> 'p' is false = not-p
>
> And to say that a proposition is whatever can be true or false amounts to saying: we call something a proposition when *in our language* we apply the calculus of truth functions to it...
>
> But this is a bad picture. It is as if one were to say "The king in chess is *the* piece that one can check." But this can mean no more than that in our game of chess we only check the king. Just as the proposition that only a *proposition* can be true or false can say no more than that we only predicate "true" and "false" of what we call a proposition. And what a proposition is is in one sense determined by the rules of sentence formation (in English for example), and in another sense by the use of the sign in the language-game... (*Investigations*, §§ 135-136)

As it is with understanding what a proposition is, it is with interpreting *any* utterance. The reason we think of the bedridden invalid as confused when he utters (ii) is that the utterance makes no sense *in that case*. If we suppose that the invalid really believes his statement when he utters it, we can only conclude that he is "out of touch with reality": he doesn't understand his own condition. The

statement makes *some* sense, to be sure: we can understand what it would mean if it were said by someone who could actually get out of bed. In the case of the invalid, that possibility is not a live one; so we can only make sense of the utterance by supposing that the invalid is confused.

The importance of the conditions of utterance to the meaning of a statement cannot be overestimated:

> Language is a characteristic part of a large group of activities -- talking, writing, travelling on a bus, meeting a man, etc. We are concentrating, not on the words 'good' or 'beautiful', which are entirely uncharacterisitic, generally just subject and predicate ('This is beautiful'), but on the occasions on which they are said... (Wittgenstein, "Lectures on Aesthetics", § 5)

As I noted above, Lewis is unlikely to object to my suggestions here. Lewis obviously thinks that context-dependency is something we cannot ignore. Moreover, he argues that the sensitivity to context is an advantage of his modal realism. He might concede that some sentences which appear to be quantified across worlds turn out to be nothing of the kind. So, he can say that modal realism should apply only to the "real" modal sentences, and not the apparent ones. Of course, that means that we have to base our formal analysis of modal sentences upon a pre-formal, ordinary-language semantics. (Otherwise, we can give up the utility of modal realism in analysing language.) For there seem to be two options to explain the apparent modal sentences which are not to be analysed with modal realism. The first is to say that, whenever an apparently modal sentence cannot easily be analysed under modal realism,

the theory does not apply to that sentence. That seems unacceptable, since it is question-begging: modal realism is used to analyse all and only sentences which are amenable to modal-realist analysis, and that ease of analysis is taken to show that modal realism is a good theory. It is the second option that Lewis seems to prefer, however. In this case, we try to understand what a sentence means before we analyse the sentence. That suggests that, to offer a modal realist analysis of a sentence, we must already know what it means. And that implies that it is the practical effect of the language -- its use -- which is most important, and not the objects to which it is supposed to be referring. In other words, if we must already have a clear idea of what a sentence means before we offer a modal realist analysis, we should ask what advantages there are in being a modal realist. We could just as well be modal agnostics: we could refuse to commit ourselves in any way to an ontology for our modal sentences.

Nevertheless, there seems to be no reason why we might not still interpret every use of modal idioms with a possible-worlds semantics; it is not inconsistent with understanding such idioms in conversation. But, as I argued above, it is not the *only* way, nor even obviously the *best* way to understand such statements. The case of (ii) seems to show as much: we can better understand the utterance as a statement of belief, or an offer of assistance, or something similar. And, such an approach allows us easily to see not only *that* such a claim is true or false, but *why*, without having to posit strange worlds which are

isolated from us. But what about the more complicated examples: (i) and (iii)? Perhaps we can understand (i) as a mere statement of belief that my life could have been different; but surely (iii) is not a claim about what my life would be like if I had the same experiences as Winston Churchill[1]. It seems that these cases, as problematic as they may be for modal realism, give us trouble however we might like to understand them. In order to fully understand such statements, then, we need one more element for a naturalistic view about language.

5.3 Modal idiom as analogy

What is it that I try to communicate when I say something like "If I had been Winston Churchill, I never would have escaped from Pretoria"? Perhaps it is a claim about what I would be like if I were Winston Churchill; but that does not seem right. For the claim seems to say something about *me*. At the same time, it cannot be that the character of Winston Churchill is beside the point: we want to say something about what might happen if someone is like *me* in Winston Churchill's position, but (obviously) therefore with some of the background that Winston Churchill had when escaping from the Boers. The best way to understand this is, as in the case of (ii), to try to understand what the utterance means given its context.

---

[1]If it were such a claim, we couldn't evaluate whether it is true or false. Nobody could know what I would be like if I had been born in a palace.

Consider a case of two interlocutors, where C tells the heroic tale of Churchill's escape from Pretoria. D says, "Wow. If I had been Churchill, I *never* would have escaped from Pretoria." What is communicated by the statement? In this case, it is obvious: D is saying something about the likelihood of his taking a tremendous and impetuous risk in order to gain some benefit. After all, in the actual case of Churchill's escape, Churchill knew he was in no danger in the Boer prison. He was well-known both at home and to the Boers, so he was not likely to be executed.[2] There is evidence, however, that Churchill's political ambitions led him to impetuous action; so, in this case, we can reasonably understand D as saying that he is not the sort of person who would try a risky operation (a fleeing prisoner of war can, of course, be shot on sight) for political fame.

Consider, however, another case. C is explaining to D that, though the idea for the escape plan had not been Churchill's, he was the only one who benefited. Churchill's co-conspirators in the escape plan never managed to evade the scrutiny of a sentry in the way that Churchill had. In this case, we can imagine D being offended by Churchill's willingness to leave his comrades behind. D might say, "Had *I* been Churchill, *I* never would have escaped from Pretoria." Here, D is making a claim about what Churchill ought to have done. How, then, are

---

[2]Indeed, Churchill was never in any real danger: the Boers realised, upon capturing him, who he was, and viewed him as something of a prize. (Manchester, *Visions of Glory*, pp 300 ff.)

we to understand the meaning of "Had I been Churchill"? In what way does it make sense to speak of my being Churchill?

What seems natural here is not to claim, somehow, that (in some world) I *am* Churchill, or vice versa; rather, it appears that (iii) functions as a comparison between the character of Churchill and the character of the speaker. What I am claiming, really, is, were I in a situation *enough like* Churchill's, I would react in a particularly different way from the way he did. Lewis invokes the closeness of worlds, to say, "Yes, and the world which is *enough like* this one is the one in which you escape from Pretoria." It is certainly true that we can use modal realism to offer an analysis; but we do not need modal realism. We can get by with a simpler, albeit primitive, notion: analogy.

When we make an analogy, we try to communicate a comparison between the analogues. If I make a statement like (iii), what I want to do is invoke some comparison between me and Winston Churchill. Moreover, I do not want the comparison between some subtly different "me", only exactly in the position of Winston Churchill. Rather, I want to pick out particular features of Churchill, his situation, etc., and point out how those compare to similar features of me, and my situation. That is, of course, not at all surprising in conversation; we understand the comparison a speaker is trying to make because of his inflection, what went before, and the like. The confusion only arises when we take the statement in isolation.

Lewis might like to argue, here, that the conversational context is not fair to his view. After all, counterfactual conditionals such as (iii) really only include "antecedent sketches": they do not make explicit what is at issue in the case at hand.

> Thus we account for such pairs of counterfactuals as Quine's
>
> *If Caesar had been in command [in Korea] he would have used the atom bomb.*
>
> versus
>
> *If Caesar had been in command he would have used catapults.*
>
> If in doubt, we ask the propounder of the counterfactual supposition that Caesar was in command whether he has in mind a modernized or an unmodernized Caesar. We thus ask him to make explicit part of his antecedent that was left implicit in his antecedent-sketch '*If Caesar had been in command...*'. (*Counterfactuals*, pp 66-7)

Lewis goes on to argue that context and conversation will allow us to clarify whether it is a modernised or unmodernised Caesar that is under consideration.[3] But Lewis's willingness to allow for the normal vagueness of conversation misses the point of the statements altogether. For there is no "fact of the matter" about Caesar's choice of weaponry in Korea, irrespective of how modern a general Caesar might be. The two statements are trying to point out features of Caesar, or the conflict in Korea, or political expediency, or something else.[4] We do not need

---

[3]This clarification might require making explicit the implicit bits of the antecedent, or it may depend solely upon conversational context; it makes no difference here. Cf. *Counterfactuals*, p 67.

[4]In this case, of course, the statements are trying to show us something about modality. My claim is that they can show us nothing

the context merely to know whether the speaker is referring to a modern or ancient Caesar; we need the context in order to evaluate the statements *at all*. Counterfactual conditionals do not appear in isolation. They are an attempt to illustrate something about the world. Even if we can know their truth conditions, or even truth values, that is beside the point.

Consider what we could use the statements about Caesar to mean. Someone might try to show that the command in Korea had not been sufficiently tough. In that case, the statement might be a comparison with Caesar's willingness to use fairly brutal methods to win wars: "If *Caesar* had been in command, *he'd* have used the atom bomb." The implication, in such a case, is that the commanders in Korea were not as good commanders (in their war) as Caesar was (in his). On the other hand, the same statement could be used to show that Caesar was needlessly brutal. Perhaps someone is arguing that Caesar had no concern for the effects of some action, as long as it ensured his immediate victory. In that case, the person might say, flippantly, "If Caesar had been in command in Korea, he would have used the *atom bomb*." What is important, here, is not whether Caesar would (really) have used the atom bomb, but what it says about Caesar, or modern commanders. We

---

just because they are not real instances of language use: nobody ever makes such a claim in isolation.

can use possible worlds to analyse such statements, but they do us no service.

Lewis might like to argue that the use of possible worlds is the easiest way to get truth-values (or, at least, truth-conditions) for counterfactual conditionals. But even if that is true, there is nothing wrong with understanding these locutions as analogical. On such a view, the "worlds" are really just analogies. In the case, (iii), where I do not escape from Pretoria, I am something like Winston Churchill, and something like the actual me. How much, and in what ways, like each of these things? Enough to be recognisably so; and, enough so that my audience can understand what (iii) means. For all we really need, in interpreting (iii), is an estimation of what the claim is supposed to communicate, and an understanding of how I am supposed to be like Churchill (in order for the claim to make sense).

Lewis would call my suggestion a form of ersatzism. And, he would say that he has shown how ersatzism cannot hold: we saw as much in chapter 2. But Lewis's criticism of ersatzism depends on the assumption that the ersatz, linguistic worlds are true and complete descriptions of some state of affairs. I see no reason to suppose that we need a complete description of a possible world. My "analogical worlds" are incomplete. But that seems to be disastrous. For it would seem that we cannot have logical entailment in these analogical worlds. But that does not seem to be a real objection: in most of the cases of counterfactual sentences, the speaker is trying to point out relevant

differences between the actual state of affairs and the described, counterfactual state of affairs. Hence, the understanding of the ersatz worlds as analogies: the state of affairs in question is *just like the real world*, except for the differences that the speaker is trying to communicate. So, analogical understanding does not need the analogical world to be a *complete* description; all it needs is *enough* understanding of the description (in the actual world) to evaluate the counterfactual conditional. In most cases, of course, we do these evaluations as a matter of course. We interpret statements about Caesar and Korea, or my Churchillian activities in South Africa, in terms of what such statements are trying to say. If we want a form, possible-world semantics for such statements, however, the option remains open to us, without committing us to the real existence of every conceivable alien realm.

My approach is modelled after the way St Thomas Aquinas understands statements about God[5]. For example, it seems wrong to claim

---

[5]St Thomas never formally outlined his doctrine of analogy. A good study is in Klubertanz, especially chapter VI. What is important about St Thomas's approach is outlined nicely there:

> When various, distinct, and independent things are conceived, known, or understood according to their perfection as beings, then they are found neither to be simply the same in this perfection nor to be so diverse that they cannot be conceived according to some kind of common perfection. The being which is said to be analogous is not the concept, for a concept cannot be predicated of a thing. It is not simply a thing. It is an intelligibility which formally as intelligibility is actually existing only in a mind. (*St. Thomas Aquinas on Analogy*, p 115)

that humans could be like God, as humans are finite and God is not (among other reasons). But, doctrine claims that humans are made in the image of God. St Thomas replies:

> Therefore if there is an agent not contained in any *genus*, its effects will still more distantly reproduce the form of the agent, not, that is, so as to participate in the likeness of the agent's form according to the same specific or generic formality, but only according to some sort of analogy; as existence is common to all. (*ST*, I 4.3)

The likeness of creation to God, then, is analogical. Similarly, we can say that the world in which I am like Winston Churchill is analogous to the actual one. Moreover, the existence of that world is also analogical: possible worlds need not be compared to the actual world as "more things of *that* sort, differing not in kind but only in what goes on in them." (*Counterfactuals*, p 85) We can, therefore, avoid the strangeness of Lewis's ontological liberalism, without giving up the philosophers' paradise that he finds in possible worlds.

We might wonder now what the ontological status of the analogical, ersatz worlds are. Clearly, since my proposal is a form of ersatzism, the worlds are not supposed to be real in just the way the actual world is. And, since it is a form of linguistic ersatzism, these are supposed to be descriptions of some kind. But, normally, an analogy requires two analogues. What are the things which are being compared with the actual worlds? I deny that the worlds are "real things", in the way that the actual world is real. But how can an unreal thing be compared with a real thing?

I answer that the analogy is between two perfectly similar things: descriptions of states of affairs. Of course, one of the analogues is the description of a real state of affairs, and the other analogue is a descripton of an unreal state of affairs. But that difference explains actuality in a much more natural way than Lewis does. For it seems perfectly natural to believe the statement, "Only actual things are real." Lewis's view makes that false. I can explain why it is true. The ersatz worlds are minimal descriptions of states of affairs. But no matter how I might describe those worlds, one of the features of those states of affairs is that they are not the case. Implicit in my utterance of examples (i) through (iii), or some statement about what Caesar might have done in Korea, or any other such counterfactual, is the knowledge that the statements describe some state of affairs which does not obtain. It makes sense, then, to suppose that (in the case of such utterances) I am making a comparison between the actual state of affairs, and some non-actual states of affairs which interest me. I use modal language to pick out those states of affairs which I take to be important to my point.

One might object, now, that I have not really said anything positive about what these analogical worlds are. I have said what they are not: they are not "real". I have also suggested that they are just descriptions of states of affairs, but that those states of affairs do not obtain. One might respond by saying that, since the analogies describe states of affairs which do not obtain, then everything one

might say about those states of affairs must be false. That objection, of course, depends upon the assumption that the existential quantifications must refer to some real, existing thing in order to be true. We discussed that assumption in chapter 4, and saw that it leads to unwelcome consequences. It seems, then, that we might be willing to accept such existential quantifications (at least in some cases) on the basis of something like a sayso semantics (cf. Woods, ch. II). My response to being a prisoner of the Boers is whatever I say it would be, just because of what I am trying to communicate. The statement would be false were it inconsistent with what I am trying to say.

The above seems to imply that, in some cases, I could truly say, "'A' and 'not-A'"? I see nothing wrong with that, as long as what I am saying depends upon the Law of Contradiction being violated -- if I am writing a fairy-tale, for instance. Of course, most of the time, all the same logical rules as apply in an actual case must apply in a counterfactual case under consideration. But that is not because of the status of the logic in various worlds. Rather, it is because we are trying to compare situations where the logic is the same. My view seems to allow for outright inconsistency; and, I say, that is a reason to think it worthwhile. In some -- very few -- cases, our natural-language conversations make such allowances. I think our understanding of modality must make similar allowances, even at the price of such ambiguity.

Lewis seems perfectly willing to accept that natural language is full of outright ambiguities; certainly, interpretation of language must be context-sensitive (see chapter 2, and *PW*, pp 40 ff.). His claim is, nevertheless, that we should use the resources of modal realism in order to interpret modal claims. The reason modal realism seems convincing to him is that it gives us the paradise of *possibilia*. But, by using analogical interpretations of modal idioms, we can have *possibilia* without the unhappy ontological costs of modal realism. Lewis might object, now, that my proposal is more liberal than is his. For modal realism is "merely quantitatively, not qualitatively, unparsimonious." (*Counterfactuals*, p 87) It posits many objects, it is true; but, those objects are all of the same kind. My proposal, on the other hand, allows different kinds of objects: the actual world, and then other things (the analogical worlds). My response is to deny that the "analogical worlds" are real things. We can quantify across worlds for the sake of formal convenience; but those "worlds" are not real things. Rather, they are descriptions of unreal things. And, though the analogical worlds are prone to the criticism that the "worldmaking language" (the analogies) cannot be complete enough for our purposes, there is nothing wrong with that. The analogies need not be complete, since our language-use has gaps (just as do our thoughts and beliefs). In actuality, we never do imagine that logical space is full: we can afford to be agnostic on the matter.

Lewis might also argue that my proposal needs to take modality as primitive. That is supposed to be a disadvantage, of course, because we want to cut down on primitive notions for the sake of analysis. But, of course, we gain nothing if we reduce the number of primitives we have by making our analysis less accurate, or by using something equally obscure. Lewis's view does both: we saw (in chapters 2 and 3) that the analysis of modal idioms seems strained, and that his view uses inscrutable notions like essences. The analogical approach I advocate takes the language-use into account for each analysis; if the analysis seems strained in any particular case, we should look for a better analysis. Similarly, we do not need essences, perfectly natural properties, or any other such mysterious things to understand modal claims. It may be true that we give up some theoretical economy by taking modality as primitive; it is no more a cost, however, than is the notion of essences.

Recall that Lewis's argument for modal realism was that the theoretical advantages it offers are worth the cost in our ontology. But we can have the advantages without the modal cost, by interpreting any reference to a world as an analogy with this world. Moreover, we can have a more natural understanding of language by interpreting many modal claims as belief statements, statements of intent, and the like. Such interpretation is natural enough: Lewis admits that interpretation needs to be context-sensitive anyway. All of that means that we *can*

have "paradise on the cheap". That seems to be a "better deal" than the one that Lewis is offering.

Bibliography

St Augustine of Hippo. *The Confessions of Saint Augustine.* trans. E. Pusey, D.D. New York: Random House, 1949.

Carnap, R. *Meaning and Necessity.* Chicago: University of Chicago, 1947.

Cassin, C.E. "Russell's Distinction between the Primary and Secondary Occurrence of Definite Descriptions" in *Essays on Bertrand Russell*, ed. E.D. Klemke. Urbana: University of Illinois Press, 1970. pp 273-84.

Davidson, D. "On the Very Idea of a Conceptual Scheme" in *Inquiries into Truth and Interpretation.* Oxford: Clarendon, 1984. pp 183-198.

----. "True to the Facts" in *Inquiries into Truth and Interpretation.* pp 37-54.

Descartes, R. *The Philosophical Writings of Descartes*, v. III, ed. Cottingham *et al.* Cambridge: Cambridge University, 1991.

Forbes, G. *The Metaphysics of Modality.* Oxford: Clarendon, 1985.

Forrest, P. "Occam's Razor and Possible Worlds" in *Monist* 65, 1982. pp 456-464.

García Márquez, G. *Love in the Time of Cholera.* trans. Edith Grossman. New York; Penguin, 1988, 1989.

Hintikka, J. "Quantifiers in Deontic Logic" in *Societas Scientiarum Fennica, Commentationes Humanarum Litterarum*, 23, no.4, 1957.

-----. "The Modes of Modality" in *Acta Philosophica Fennica* 16, 1963, pp 65-79. Reprinted in Loux, pp 65-79. Page references are to the reprint.

James, W. "The Will to Believe" in *The Writings of William James: A Comprehensive Edition*, ed. John J. McDermott. Chicago: Univiersity of Chicago, 1977. pp 717-735.

Kanger, S. *Provability in Logic.* Stockholm: Almquist and Wicksell, 1957.

Klubertanz, G., S.J. *St Thomas Aquinas on Analogy: A Textual Analysis and Systematic Synthesis.* Chicago: Loyola University, 1960.

Kripke, S. "A Completeness Theorem in Modal Logic" in *Journal of Symbolic Logic*, 24, 1959. pp 1-14.

-----. "Semantical Considerations on Modal Logic" in *Acta Philosophica Fennica* 16, 1963, pp 83-94. Reprinted in Linsky, pp 63-72. Page references are to the reprint.

Leibniz, G.W., *Monadology* in *Philosophical Writings*, ed. G.H.R. Parkinson. trans. M. Morris and G.H.R. Parkinson. London: J.M. Dent & Sons, 1973. pp 179-194.

-----. *Necessary and Contingent Truths* in *Philosophical Writings*, pp 96-105.

Lewis, C.I., and C. Langford. *Symbolic Logic.* New York: Dover, 1932, 1951.

Lewis, D. *Counterfactuals.* Cambridge, Mass.: Harvard University, 1973.

-----. "Counterpart Theory and Quantified Modal Logic" in *Journal of Philosophy* 65 (1968), pp 113-126. Reprinted with postscripts in Lewis, *Philosophical Papers*, v I. New York: Oxford University, 1983. pp 26-46. All citations are from the reprint, with postscripts.

-----. *On the Plurality of Worlds.* Oxford: Blackwell, 1986.

Linsky, L., ed. *Reference and Modality.* London: Oxford University, 1971.

Loux, M.J. "Modality and Metaphysics" in Loux. pp 15-64.

Loux, M.J., ed. *The Possible and the Actual: Readings in the Metaphysics of Modality.* Ithaca: Cornell University, 1979.

Manchester, W. *The Last Lion: Winston Spencer Churchill; Visions of Glory, 1874-1932.* New York: Laurel, 1983.

Montague, R. "Logical Necessity, Physical Necessity, Ethics, and Quantifiers" in *Inquiry* 3, 1960. pp 259-269.

Plantinga, A. *The Nature of Necessity.* Oxford: Clarendon, 1974.

-----. "Transworld Identity or Worldbound Individuals?" in *Logic and Ontology*, ed. Milton Munitz. New York: New York University, 1973; reprinted in Loux. pp 146-165.

Quine, W.V. *From a Logical Point of View*, 2d ed., rev. Cambridge, Mass.: Harvard University, 1964.

Quine, W.V. *Pursuit of Truth.* Revised ed. Cambridge, Mass.: Harvard University, 1992.

-----. "Reference and Modality" in *From a Logical Point of View.* pp 139-159.

Russell, B. "On Denoting" in *Mind*, New Series, XIV, Oct. 1905. pp 479-93. Reprinted in *The Collected Papers of Bertrand Russell* v 4, *Foundations of Logic, 1903-05*, ed A. Urquhart and A.C. Lewis. New York: Routledge and The Bertrand Russell Editorial Project, 1994. pp 414-27. All citations from the reprinted edition.

Russell, B. *Principles of Mathematics.* 2d ed. New York: Norton, public domain.

Schlesinger, G.N. "Possible Worlds and the Mystery of Existence" in *Ratio* 26, 1984. pp 1-17.

St Thomas Aquinas. *Summa Theologica. The "Summa Theologica" of St Thomas Aquinas.* trans. Fathers of the English Dominican Province. Public domain.

Whitehead, A.N. & B. Russell. *Principia Mathematica.* v.1. Cambridge: Cambridge University, 1910.

Woods, J. *The Logic of Fiction: A philosophical sounding of deviant logic.* The Hague: Mouton, 1974

Wittgenstein, L. "Lectures on Aesthetics" in *Lectures and Conversations on Aesthetics, Psychology and Religious Belief.* From notes by Y. Smythies *et al.*, ed. Cyril Barrett. Berkeley: University of California, N.D.

-----. *Philosophical Investigations.* 2d ed. ed. G.E.M. Anscombe and R. Rhees. trans. G.E.M. Anscombe. New York: Blackwell & Mott, 1958.