NOTICE

The quality of this microform is heavily dependent upon the quality of the original thesis submitted for microfilming. Every effort has been made to ensure the highest quality of reproduction possible.

If pages are missing, contact the university which granted the degree.

Some pages may have indistinct print especially if the original pages were typed with a poor typewriter ribbon or if the university sent us an inferior photocopy.

Reproduction in full or in part of this microform is governed by the Canadian Copyright Act, R.S.C. 1970, c. C-30, and subsequent amendments.

AVIS

La qualité de cette microforme dépend grandement de la qualité de la thèse soumise au microfilmage. Nous avons tout fait pour assurer une qualité supérieure de reproduction.

S'il manque des pages, veuillez communiquer avec l'université qui a conféré le grade.

La qualité d'impression de certaines pages peut laisser à désirer, surtout si les pages originales ont été dactylographiées à l'aide d'un ruban usé ou si l'université nous a fait parvenir une photocopie de qualité inférieure.

La reproduction, même partielle, de cette microforme est soumise à la Loi canadienne sur le droit d'auteur, SRC 1970, c. C-30, et ses amendements subséquents.

THE UNIVERSITY OF ALBERTA

THE MAGIC OF DAVID LEWIS

by

Tilman Lichter 

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of Master of Arts.

Department of Philosophy

Edmonton, Alberta
Spring, 1990

ISBN 0-315-60128-0

THE UNIVERSITY OF ALBERTA

RELEASE FORM

Name of Author:	Tilman Lichter
Title of Thesis:	The Magic of David Lewis
Degree for which thesis was presented:	Master of Arts
Year this degree granted:	Spring, 1990

Permission is hereby granted to the **University of Alberta Library** to reproduce single copies of this thesis and to lend or sell such copies for private, scholarly or scientific research purposes only.

The author reserves other publication rights, and neither the thesis nor extensive extracts from it may be printed or otherwise reproduced without the author's written permission.

(Signed) .....

Permanent Address:

9614 78 Ave

Edmonton, AB

T6C 0P3

Date April 25 ....19 90.

THE UNIVERSITY OF ALBERTA

FACULTY OF GRADUATE STUDIES AND RESEARCH

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research, for acceptence, a thesis entitled THE MAGIC OF DAVID LEWIS submitted by Tilman Lichter in partial fulfillment of the requirements for the degree of Master of Arts.

Mohan Matthen ..... Mohan Matthen

Supervisor

Francis Jeffry Pelletier .....

Martin M. Tweedale .....

R. Goebel .....

Date April 6th 1990

Abstract

In his book, **On the Plurality of Worlds**, David Lewis defends a theory of modality he calls 'modal realism'. He argues that, just as the utility of set theory in mathematics gives mathematicians good reason to believe in a plurality of sets, so the fruitfulness of presupposing a plurality of worlds for modal applications should give philosophers good reason to believe in a vastly expanded ontology of possibilia. In order to gain 'honest title' to the theoretical benefits which talk about possibilia can make available to us, Lewis argues, we should accept this kind of talk as the literal truth. This entails accepting the existence of a vast realm of worlds just as real and concrete as our own, populating 'logical space' in complete causal and spatiotemporal isolation from us and from each other. There are so many worlds in this 'paradise for philosophers', as Lewis calls it, that 'absolutely every way that a world could possibly be is a way that some world is'.

This controversial proposal has become the focus of critical debate by a number of modal theorists less inclined to embrace exotic ontologies as long as they remain unconvinced of the actual benefits of such a move. Lewis, by contrast, claims that the competing opinions all exhibit serious shortcomings, and he devotes a portion of his book to a careful exposition of these problems.

What remains to be examined, and what therefore constitutes the core of this thesis, is whether Lewis own modal theory can in fact deliver where his rivals fail. The main difficulties to be considered in this context are:

1. the problem of reference
2. the problem of analyzing relations
3. the problem of primitive modality

My thesis concludes that, although David Lewis's theory does exhibit improved unity and economy, it cannot overcome the obstacles which had seemed most damaging to his competitors positions. Thus, the doubt which, by his own admission, troubles Lewis the most seems well justified: the benefits of his theory are not worth the cost. Modal realism suffers from the same debilitating defects which prevent competing theories from providing an analysis of modality. The benefits which modal realism had promised to deliver can be had more cheaply elsewhere.

TABLE OF CONTENTS

Introduction 1

Section 1: Modal Realism 7

1.1 Possible Worlds 7
1.1.1 What Makes a Lewis-World a World? 14
1.1.2 Causal Isolation 18
1.1.3 Concrete or Abstract? 20
1.1.4 Actuality 24

1.2 Properties 26
1.2.1 Abundant or Sparse? 28

1.3 Propositions 30

Section 2: The Ersatzist Challenge 32

2.1 Linguistic Ersatzism 33
2.1.1 Primitive Modality 34
2.1.2 Descriptive Power 37
2.1.3 Indiscernible Individuals 43
2.1.4 Summary 43

2.2 Pictorial Ersatzism 45
2.2.1 Primitive Modality Revisited 47
2.2.2 Conflation of Indiscernibles 48
2.2.3 The Collapse of Pictorial Ersatzism 49

2.3 Magical Ersatzism 52
2.3.1 External or Internal? 54

2.4 Summary 58

Section 3: The Magic of David Lewis 59

3.1 The Indeterminacy of Reference 60
3.1.1 The Model-Theoretic Component 61
3.1.2 The Impact on Model-Theoretic Semantics 63
3.1.3 Meaning Just Ain't In The Head 67
3.1.4 Does Lewis Have No Models? 73

3.2. Magical Relations 79
3.2.1. Truth Value Selection 81

3.3 The Principle of Recombination 90

3.3.1 The Worlds According to Hume 91
3.3.2 The Problem With The Intended Interpretation. 102

Bibliography 111

INTRODUCTION

> Incidentally, recognition of multiple worlds or true versions suggests innocuous interpretations of necessity and possibility. A statement is necessary in a universe of worlds or true versions if true in all, necessarily false if true in none, and contingent or possible if true in some... But such an account will hardly satisfy an avid advocate of possible worlds any more than spring water will satisfy an alcoholic. Nelson Goodman[1]

> Necessarily all swans are birds iff, for any world W, quantifying only over parts of W, all swans are birds. More simply: iff all swans, no matter what world they are part of, are birds. The other modalities follow suit. What is impossible is the case at no worlds; what is contingent is the case at some but not at others. David Lewis[2]

When Nelson Goodman wrote the passage above, he evidently did not understand the true cravings of avid advocates of possible worlds, or else he misjudged their ability to come to their senses and see, like the reformed alcoholic, that it is indeed spring water which is the eminently, and the only lastingly satisfying, substance.

David Lewis, possible world devotee par excellence, is a good case in point. At least on the face of it, spring water seems to do for him what Goodman thought possible to achieve only with much harder stuff. For when Goodman dismisses possible world talk, saying:

---

1. Nelson Goodman, 1978, **Ways of Worldmaking**, Indianapolis, Hackett Publishing Co., p.120.
2. David Lewis, 1986, **OPW**, Oxford, Basil Blackwell Ltd., p.7.

> Let it be clear that the question here is not of the possible worlds that many of my contemporaries, especially those near Disneyland, are busy making and manipulating. We are not speaking in terms of multiple possible alternatives to a single actual world, but of multiple actual worlds.[3]

Lewis can cheerfully support this:

> [E]very world is actual at itself, and thereby all worlds are on a par.[4]

Yet Goodman writes his **Ways of Worldmaking** in order to argue that our universe consists of nothing but the ways we have of describing it, rather than consisting of some neutral world or worlds underlying those descriptions. As he sees it, we may be able to compare our conflicting descriptions, or versions, of what we take to be the 'real world' with each other, but we "cannot test a version by comparing it with a world undescribed, undepicted, unperceived,... and while the underlying world... need not be denied to those who love it, it is perhaps on the whole a world well lost."[5] He is convinced that such a world would be "a world without kinds, or order or motion or rest or pattern - a world not worth fighting for or against."[6]

---

3. Goodman, 1978, p.2.
4. Lewis, 1986, p.93.
5. Goodman, 1978, p.4.
6. Goodman, 1978, p.20.

David Lewis, by contrast, declares: "our world is but one among many",[7] and goes on to clarify in direct opposition to Goodman's claim:

> The worlds are not of our own making... We make languages and concepts and descriptions and imaginary representations that apply to worlds. We make stipulations that select some worlds rather than others for our attention. Some of us even make assertions to the effect that other worlds exist. But none of these things we make are the worlds themselves.[8]

To be sure, Lewis is aware of some dissenters who would assert that, quite to the contrary, we often do make descriptions or pictures of whole worlds, and that these descriptions and pictures are the worlds themselves. He devotes a lengthy discussion under the title **The Ersatzist Programme** to a detailed refutation of their claims. But despite some interesting similarities between Goodman's and Lewis's accounts, Lewis does not include **The Ways of Worldmaking** in either his bibliography or his critique of linguistic ersatzism. Correctly so, in my opinion, because in **On the Plurality of Worlds (OPW)** it is his intention to present a defense of what he calls "modal realism", beginning with the fundamental hypothesis that whatever and however many worlds there are, exist simpliciter, independent of what anyone might think or say about them. He claims that his hypothesis is necessary for a correct analysis of

---

7. Lewis, 1986, p.2.
8. Lewis, 1986, p.3.

modality. He admits that, although he has good reason for this claim, he does not have conclusive reason for it. Consequently, what interests Lewis most is any rival claim making the benefits of his theory not worth the cost for the simple reason that they can be had more cheaply elsewhere.

It will be the aim of this thesis to examine in some detail the sort of benefits we might expect to reap as a result of embracing the modal realist belief in a plurality of worlds, and why other theories fail to deliver them. However, we will also have to ask if Lewis is correct in asserting that his own theory has the power to make the promised benefits available to us.

In order to facilitate discussion of these matters, section 1 will be devoted to a careful exposition of some of the relevant aspects of Lewis's modal realism. Section 2 will present Lewis's criticism of the 'ersatzist programme' as well as a critical evaluation of his arguments in light of his own theoretical commitments. I will try to show that Lewis fails to supply a generally acceptable explication of the kind of 'modal facts' our strongest intuitions lead us to expect behind natural language statements containing the terms 'possibly' and 'necessary'. Despite the fact that it is his avowed purpose to render natural language talk about possibility

more perspicuous and amenable to truth-theoretic analysis, his metaphysical 'worldmaking' leads him in the end to a postulated reality which allows him to assign truth values to a great majority of modal sentences, while at the same time requiring us to radically change some of our most basic intuitive convictions about the nature and purpose of modal talk. As Alvin Plantinga has remarked, he is comparable to a

> hyper-liberal theologian who insists that on his theory there is such a person as God all right, even though there are no supernatural beings: for the word 'God', as he uses it, he says, denotes the evolutionary-historical process (or perhaps "the forces not ourselves that make for goodness"). Suppose such a theologian goes on to model the rest of what theists ordinarily say in nonsupernatural beings: he doesn't share their belief that there is such a person as God, even though on his theory the words 'there is such a person as God' express a truth. His opinion differs from that of the plainspeaking atheist only by virtue of being less plainly spoken.[9]

Lewis succeeds in demonstrating the structural isomorphisms between his postulated ontological reality and the sentences of modal language only at the cost of making the relation between them by virtue of which modal statements would become capable of referring to such a reality in the first place either primitive or else magical.

---

9. Alvin Plantinga, 1987, **Two Concepts of Modality**, Philosophical Perspectives, 1, Atascadero, Ridgeview Publishing Co., pp.225-226.

Section 3 will try to investigate the reasons for this result from the point of view of some recent considerations by Hilary Putnam on the ultimate viability of model-theoretic semantics. The claim to be substantiated will be that the choice between Lewis's modal realism and the modal theories of other philosophers is not a choice between a theory which offers a proper analysis of modality and competing systems which are either inconsistent or else accept modality as primitive. Rather, Lewis fails to achieve an analysis of modality for much the same reasons his competitors do. Modality remains a mystery for Lewis as much as for linguistic or so-called magical ersatzers. The choice we have is between a system with an extremely outsized and exotic ontology, and some others who do not strain our credulity quite as much. In the beginning of OPW Lewis said:

> Maybe - and this is the doubt that most interests me - the benefits are not worth the cost, because they can be had more cheaply elsewhere.[10]

I think his doubt was indeed well founded. The benefits are not worth the cost. We would be well advised to buy the less pricey modal theories on offer elsewhere.

---

10. Lewis, 1986, p.5.

SECTION 1: MODAL REALISM

1.1 Possible Worlds

In OPW, David Lewis is concerned with the task of making systematic philosophizing easier for philosophers. The problem he singles out for his attention is the peculiar human propensity to imagine and talk about the many ways in which things, events, or states of affairs might have been other than they actually turned out to be. Modal talk, talk about what is possible and what is necessary, puzzles Lewis as it has puzzled philosophers for millennia.

But Lewis is not primarily interested in a metalogical investigation of modal phenomena.

> Metalogical results, by themselves, answer no questions about the logic of modality. They give us conditional answers only: if modal operators can be correctly analysed in so-and-so way, then they obey so-and-so system of logic. We must consider whether they may indeed be so analysed; and then we are doing metaphysics, not mathematics.[1]

In other words, given that logicians have had some success in constructing formal systems of modal logic capable of mirroring many of our intuitive notions about modality, how are we to decide between competing systems unless we have some idea how our operators should be interpreted? Given, for instance, that we interpret them as restricted quantifiers over a set of entities, how can

1. Lewis, 1986, p.17

we settle the status of controversial axioms in quantified modal logic without knowing what sort of entities we are quantifying over? As Lewis observes, we need a substantive analysis of modality to answer these questions. And he claims that his own theory, modal realism, provides us with all the necessary tools.

What then is modal realism?

As elaborated by David Lewis, the theory of modal realism is without much doubt a most thoroughgoing application of the Quinean dictum 'to be is to be the value of a variable' to metaphysical investigation. Guided by the apparent isomorphic structure of some systems of quantified modal logic and parts of natural language talk about possibility and necessity, Lewis proposes to take the existence claims of the quantifiers in formulas of quantified modal logic as seriously as we would be inclined to take them in non-modal quantificational systems. It is this suggestion which served to make him the target of what he has called "the incredulous stare"[2] from many of his colleagues. But it is also a move whose very boldness and radical departure from conventional views yielded a theory of great internal consistency and analytic clarity.

To put it bluntly, Lewis tells us to accept not only the existence of the very world we live in, but also the

2. Lewis, 1986, p.133

existence of countless other 'possible worlds', ontologically on par with ours. Just as our actual world might be seen as some big physical object, comprising everything which exists "at some distance and direction from here, or at some time before or after or simultaneous with now"[3] , so likewise there are countless other worlds in existence, albeit in complete spatiotemporal isolation from ours. These worlds are not our creation. They exist simpliciter in whatever sense our own universe can be said to exist. The existence of our world and the existence of those other worlds differs not in the manner of their existing, but only in their being spatiotemporally isolated from each other.

Moreover, there are so many of these independently existing worlds 'out there' that "absolutely every way that a world could possibly be is a way that some world is"[4] . And one of these countless possible worlds is the world we ourselves live in. It is this vast realm of possible entities existing in spatiotemporally isolated possible worlds, including our own, that the quantifiers of our modal language range over.

> Conversely, since it is safe to say that no world is any way that a world could not possibly be, whenever there is some world at which such-and-such is the case, then it might be that such-and-such is the case. So modality turns into quantification: possibly there are blue swans iff, for some world W, at W there are blue swans.[5]

---

3. Lewis, 1986, p.1
4. Lewis, 1986, p.2
5. Lewis, 1986, p.5

To be sure one great advantage of seeing things Lewis's way is the increase in explanatory power it provides us. Quite apart from notorious problems of satisfying modal formulas having to do with the contingency of existence when using boxes and diamonds as sentential operators, it is an undisputed fact that natural language is capable of a range of modal expressions far beyond what can be mirrored in standard modal logic unless we keep adding further and further extensions which often strike us as uncomfortably ad hoc. (Take, for instance, G. Forbes's stipulation of an "outer domain" and an "inner domain" of things existing at a world W in order to solve satisfaction in absentia problems for modal formulas.[7]) For Lewis, these natural language resources are indicative of the need for a new theory of modality capable of formalizing them in a way which preserves at least some of the structural characteristics of the original modal expressions. Two examples Lewis uses may serve to illustrate this point.

a. Numerical Quantification

'Possibly pigs fly' might be satisfied in a number of quite different ways, if three different possible individuals are pigs. Friends of boxes and diamonds would have to resort to some sort of numerical quantification: pigs fly_1, pigs fly_2, etc., thus requiring the

7. Graeme Forbes, 1982, **Canonical Counterpart Theory**, Analysis 42, pp.33-37

introduction of new entities to represent the different ways in which one and the same formula can be satisfied. Modal realism, on the other hand, possesses the resources to quantify over these 'ways' without any proliferation of entities. All that is required are the possible worlds, the individuals themselves, or set theoretic constructions of these.

b. Supervenience

The case of supervenience is perhaps the most illustrative of the particular insufficiency of boxes and diamonds, simply because supervenience is such an obviously modal issue. Intuitively, when we want to express the idea behind the term supervenience, we say something like: supervenience means that there could be no difference of the one sort without a difference of the other sort, as Lewis puts it. What interests us in this sentence is the modality contained in it. Lewis calls that the "one simple, easy, useful idea"[8] of supervenience.

Boxes and diamonds, however, restrict our talk of supervenience to particular worlds. Taking the diamond to be a sentential operator modifying sentences, we reformulate our original ordinary language statement about supervenience to read:

---

8. Lewis, 1986, p.14

> It is not the case that there is some world W such that, at W, two things have a difference of the one sort, but not of the other.

This may be a useful way of expressing facts of supervenience concerning things existing in the same world, but as before, the one feature of primitive modal logic most detrimental to a full treatment of the facts of supervenience is its inability to overcome its worldboundness. This becomes most clearly apparent with questions concerning the supervenience of laws.

Suppose we were to ask if any two worlds could possibly differ in their physical laws without at the same time differing in their distribution of local qualitative character. Reading 'could' as a diamond we would have:

> It is not the case that there exists a world W wherein two worlds differ in their laws without differing in their distribution of local qualitative character.

Even if we are somewhat less accomplished logicians, this will strike us as trivial. We are quantifying over the one world W, and there is no one world in which two worlds do anything. The trouble with boxes and diamonds, time and time again, is that they are restricted quantifiers relative to worlds.

Modal realism, by contrast, can offer us a way to unrestrict our quantifiers, so that they can range over the entire scope of possibilia, over all the things in

all the worlds. The foregoing examples seem to drive home the same point: in order to capture the richness of ordinary language modal talk we need possibilia to quantify over, "other-worldly things", as Lewis calls them.

The question remaining, however, is whether or not we have to commit ourselves to the sort of outsized ontology Lewis proposes in order to achieve our goal. Most of section 2 will be devoted to a discussion of this question. But we cannot even begin comparing the competing analyses of metaphysical necessities underlying modality without a clear idea of what conception of metaphysical reality the competing claims are based on. Granting Lewis his first premiss concerning the need for a modal language with the power to quantify restrictedly as well as unrestrictedly over a domain of entities belonging to different 'worlds', we must now investigate and analyze the key concepts of Lewis's theory. Only after we understand the ways in which his definition of these modal (and non-modal) concepts differs importantly from those of other theories of modality will we be in a position to judge if 'the price is right'. So let us turn now to the substantive part of Lewis's theory of modality and ask what has to be one of the more obvious questions:

1.1.1 What Makes a Lewis-World a World?

How are worlds internally unified? What individuates them? Apart from the fact that their mode of existence is presumably the same as any other world, Lewis states that a world is the maximal mereological sum of all the possible individuals that are parts of it. But, as he himself admits, mereology alone cannot provide us with the unifying glue for entire worlds. Obviously we don't want just any mereological sum of possible individuals to gain possible world status. Some such sums will only be parts of bigger worlds, such as for instance the mereological fusion of my desk, my chair, my keyboard, and myself. Other sums might comprise parts of quite different possible worlds, such as the mereological sum of all possible bald humans, or all possible desks.

What we need is some specifiable relation holding between the parts of one world which makes these parts worldmates, and which is restricted to the parts of one world. It is the aforementioned spatiotemporal relatedness of parts of one world which Lewis suggests to consider our 'worldmaking glue'. Thus, any maximal mereological sum of entities which are spatiotemporally related to each other constitutes a possible world. At the very least, this would give us a way to distinguish our own world from all the ones which are fairly similar to ours. But it wouldn't help much if we believed in the

possibility of worlds with a 'worldmaking glue' so alien to ours as to make it impossible to characterize it as spatiotemporal. If we indeed believe in such alien worlds, and there seems to be no prima facie reason not to, we will need an alternative definition of what it is that makes the mereological sums of parts of worlds into distinct and complete worlds. Lewis proposes to modify the initial suggestion, and to call anything a proper worldmaking relation that can be considered either strictly or analogically spatiotemporal, provided the relation is

a. natural, i.e. non-disjunctive.

b. pervasive, i.e. if it connects a chain of things with one another, then also it connects them directly with each other.

c. discriminating, i.e. it has to be at least possible for the relation to interrelate a great many entities by assigning each of them a distinct place in the overall structure.

d. external, i.e. the relation must not supervene on the intrinsic properties of the relata.

Initially it might be hoped that we can get away with a much simpler characterization of what it is that individuates worlds. Why not just say that what unifies a world is the external interrelatedness of its parts, never mind what sort of interrelatedness? On this view, all the parts of one world bear some such relation to each other, but no part of one world bears this relation to any part of another world. Thus the rather convoluted

stipulation of 'analogical spatiotemporal relatedness' might seem avoidable.

But, as Lewis points out, it is very difficult in practice to come up with an external, natural relation apart from strictly or analogically spatiotemporal relations which performs the task of interrelating all and only the worldmates. In fact, it seems easier to envision relations which would prove the attempt at simplification doomed to failure.

Imagine, suggests Lewis, that positive and negative charge were not the intrinsic properties of elementary particles we commonly take them to be. Let there be a possible world where the properties of like-chargedness and opposite-chargedness of elementary particles are instead natural external relations (according to the definition of 'natural' and 'external' above) between particles by definition. In such a world, an electron and a positron might be intrinsically identical with respect to their individual properties. Their worldmate relation would not supervene on their respective intrinsic character. But despite being natural, external, and pervasive, such a relation would not be discriminating. It would, at most, succeed in individuating three distinct groups of otherwise undifferentiated particles, namely neutral, negative, and positive particles. So its seems that this possible relation would qualify as

natural and external, but would not be analogically spatiotemporal.

Unfortunately, there also seems to be no reason to suppose that it would be impossible for parts of different worlds to be related to each other in this way, thus foiling our attempt to consider any natural, external relation whatever sufficient for the individuation of worlds from each other. Analogically spatiotemporal relations, on balance, seem to be as far as we can go in giving a comprehensive characterization of proper worldmate relation.

1.1.2 Causal Isolation

Given that worlds are spatiotemporally isolated from each other, what about causation between worlds? Must we not identify causal isolation as another individuating principle between worlds and unpack this notion theoretically?

Lewis says no. On his account of modal realism, causal isolation follows automatically from a counterfactual account of causation. To understand how this might be so, let us look at some cases. Assume we have a world W where cause C brings about event E. Events C and E both occur in the same world W and are distinct events, and it is the case at W that if C had not occurred, E would not have occurred either. Lewis's counterfactual analysis of causation amounts to saying that "at the closest worlds to W at which C does not occur, E does not occur either."[9]

It is not difficult to demonstrate that, under this particular interpretation of causation, transworld causation becomes logically implausible. Suppose event C occurred at world W(c), and event E occurred at world W(e). The counterfactual stating that if C had not occurred, E would not have occurred either, ought to hold in both these worlds if transworld causation is to make sense to us. Thus we would want to say:

---

9. Lewis, 1986, p.78

> At the closest worlds to W(c) at which C does not occur, E does not occur either.

And likewise:

> At the closest worlds to W(e) at which C does not occur, E does not occur either.

But both of these sentences fail to capture what we had wanted to capture about the fact of transworld causation. The first does not address the issue. For if we want to investigate the possibility of events in one world causing events in another, worlds close to W(c) may be informative on questions concerning the causative event C, but they do not elucidate the facts - if there be any - of transworld causation. For answers to that we have to look at worlds close to W(e).

The second sentence does not help us either. The most obvious candidate for the world closest to W(e) at which C does not occur is clearly W(e) itself!

Although Lewis goes on to develop his point in some detail, I don't want to discuss all its aspects here. Suffice it to say that, given a counterfactual analysis of causation, his account of the issue seems plausible. The purpose of this paper is to investigate the consequence of Lewis's metaphysical view of possibilia - understood as a comprehensive theoretical unity - for our developing concept of modality. In this regard, causal isolation, being one of the demarcating features of

Lewis-worlds, might be accepted without argument, if need be as a primitive principle. The preceding short exposition of Lewis's views about transworld causation is given solely to demonstrate the elegance with which he is able to make the different aspects of his theory cohere and interrelate.

1.1.3 Concrete or Abstract?

a. The Way of Example

According to this method, we understand the difference between 'abstract' and 'concrete' by giving lists of things which are paradigmatically one or the other. Thus, donkeys, puddles, protons, and stars, to use Lewis's own examples, are concrete entities, whereas, at least on some accounts, numbers are abstract entities.

Although Lewis confesses not to know exactly where to put the dividing line between the donkey-like things and the number-like things, on balance he feels more inclined to consider worlds donkey-like than number-like.

"But I know not why."[10]

The upshot of trying to settle the question by 'way of example' seems to be that, no matter how we may choose to define the respective ontological status of donkeys or numbers vis a vis each other, worlds should turn out to be on the donkey side rather than on the number side.

---

10. Lewis, 1986, p.83

b. The Way of Conflation

According to this way of explaining the difference between abstract and concrete, the difference is just that between individuals and sets, or between particulars and universals. For Lewis, this is another useful way of explaining the status of worlds according to his view. Worlds, he says, are particulars rather than universals, they are individuals rather than sets. Worlds, then, are concrete rather than abstract given that the 'way of conflation' accurately reflects the distinction between abstract and concrete.

c. The Negative Way

Here the distinction between abstract and concrete is drawn in terms of properties which concrete entities are thought to possess essentially, whereas abstract entities lack these very properties. Consequently, abstract entities lack spatiotemporal location, they do not enter into causal interactions, and they are never indiscernible from one another.

How do Lewis-worlds fit into such a scheme of things? By definition, they are spatiotemporally isolated from each other, and there is no transworld causation. The inability to enter into causal interactions is an essential characteristic of abstract entities according

to the 'negative way'. Does this mean that Lewis-worlds are abstract? If we answer in the affirmative, we will have to concede that this makes our own world abstract as well. But since we would presumably want to consider ourselves concrete rather than abstract, this would necessitate construing worlds as abstract wholes composed of concrete particulars. Making the distinction depend on the particular world from which one considers the question will not help either. To say that, for us, our own world is concrete, and all other worlds are abstract, whereas for the inhabitants of another world, their world is concrete and ours is abstract, fails to account for the kind of distinction the two terms were supposed to capture. Presumably, calling something abstract rather than concrete is meant to point to a fundamental difference in kind between the two kinds of entities so denoted.

The negative way, then, does not seem to provide us with an unproblematic distinction applicable to possible worlds. Lewis suggests to read the way of negation charitably, allowing the concreteness of parts of worlds to infect the whole. Interpreted in this way, worlds are concrete since their parts are spatiotemporally related to one another and interact causally with each other.

d. The Way of Abstraction

This may be the most commonly held view of the distinction between abstract and concrete. To be abstract means simply to be somehow an abstraction from something concrete. An abstract entity is a concrete entity from which we have 'abstracted' its specificity.

For Lewis, this makes possible worlds clearly concrete entities, since they do not lack in specificity, and there is nothing concrete that they might have been abstracted from.

On balance, then, Lewis seems justified in calling his worlds concrete. His misgivings in the matter appear to stem mostly from the somewhat ambiguous status of sets and universals, which come out on different sides of the equation depending on the criterion used. Since our question concerned the status of worlds, and since our own world should count as the paradigmatic example of what we understand a world to be, his doubts seem a little odd. After all, universals, tropes, sets, etc., are topics of discussion in this world, and their ontological status is very much a matter of philosophical concern, without therefore making it problematic to call this world we are part of 'concrete' rather than 'abstract'.

In summary, Lewis would like us to accord to possible worlds whatever status we would be prepared to accord to the world we live in. And that I think, pace the good bishop Berkeley, would be the status of concreteness.

1.1.4 Actuality

Given Lewis's acceptance of a plurality of worlds, each as 'real' or 'concrete' as our own, it seems as if we should then consider them just as 'actual' as ours. For what would the meaning of 'actual' consist in, if not in something's being 'real' and 'concrete'? And isn't it just analytically true to say that 'everything is actual'?

Lewis denies the supposed analyticity of the foregoing sentence, and provides a new definition of the term 'actual', one which will be more in line with the constraints of his theory of modality.

According to modal realism, actuality is a relative term. It is relative to particular worlds. The term 'actual' is synonymous with the term 'this-worldly'. When we use the term here in our world, it applies to our world and what is in it. When someone uses it on some other, possible world, it applies to that possible world and its constituents. Lewis:

> I suggest that 'actual' and its cognates should be analyzed as indexical terms: terms whose reference

> varies, depending on the relevant features of the context of utterance. The relevant feature of context, for the term 'actual', is the world at which a given utterance occurs. According to the indexical analysis I propose, 'actual'(in its primary sense) refers at any world W to the world W.[11]

Lewis's reason for considering 'actual' to be a relative term is twofold. In order to understand him, let us suppose we considered 'actual' to be absolute. With respect to worlds, then, actuality would be a property belonging simpliciter to one specially distinguished world to the exclusion of all others. Given the thesis of the plurality of worlds, such an understanding of actuality would still not dispense with our relative distinctions between worlds. It would still be true that there is only one single world which is ours, namely the one we happen to be part of. If actuality were the absolute property we previously supposed it to be, we would need some independent reason to convince us why, of all the possible worlds, it should just have turned out that our world ended up the one which had this elusive property. In terms of probability, the chances that our world should turn out to be the actual one were incredibly small. How is it then that we are so sure about the fact that our world is indeed the actual one?

Could it be that our acquaintance with our own actuality is itself a primitive fact? Maybe we are

---

11. Lewis, 1986, pp. 92-93

acquainted with it the way we are acquainted with our thoughts? But then, what about other people in other possible worlds? Could they not be just as immediately acquainted with their own thoughts, and thus with their own actuality?

And secondly, it is a contingent, and not a necessary matter which of all the worlds is actual. It surely could have turned out otherwise. But saying that just means that actuality is a matter which varies from world to world. That is what it means for something to be contingent. At one world, one world is actual, at another, another. We cannot square this account with the notion of absolute actuality. Given the plurality of worlds, ontologically on par with each other, we seem not to have any other options but to give the indexical analysis of actuality which Lewis proposes.

1.2 Properties

Now that we have had a first view of the kind of reality a modal realist believes himself to be living in, we are in a position to appreciate one of the first fruits such a position allows Lewis to harvest: a clear definition of the nature and role of properties.

If we believe in possible worlds, and in parts of possible worlds, and if we are prepared to accept set-theoretic constructions whose members are the things

these worlds consist of, then Lewis has a tempting proposal for us: why not take a property to be simply the set of all its instances? The property of being a donkey thus becomes defined as the set of all donkeys in all possible worlds.

There is, to be sure, a traditional problem with construing properties as sets of all their instances. The problem is that different properties might just happen to be coextensive. In our world, for instance, the property of having a heart, and the property of having a kidney are coextensive, yet are two quite different properties. A set-theoretical construal of properties would not be able to preserve this difference.

But modal realism can solve this problem. It is still the case that the property of cordateness and the property of renateness are coextensive in our world. But under a modal realist view of things, this no longer makes the two properties coextensive simpliciter. If we consider all the instances - and according to modal realism this includes all the instances in other possible worlds such as possible creatures with hearts but no kidneys, and possible creatures with kidneys but no hearts, we never encounter accidentally coextensive properties which might not have been coextensive.

But if properties are sets of this kind, they must have their members essentially. How could it ever be

contingent if a thing does or doesn't have a property, if the full membership of a set never varies?

Contingency, answers Lewis, is relative to worlds. Although the full membership of a property-set never varies from world to world, the subset which is instantiated at a particular world is a contingent matter.

Thus it is contingently true at our world that there are no flying pigs, because the property has no instances here, although there are other worlds where it does. To say that a property is contingent says nothing about the overall membership of that particular set. It merely states that not every possible world contains instances of it. Likewise, I am a member once and for all in both the set of cordateness and the set of renateness, and in our world, any creature which belongs to one of these two sets also belongs to the other. What makes this situation contingent is the fact that it is not that way in every world.

1.2.1 Abundant or Sparse?

If properties are the sets of their instances, then they must be incredibly abundant. Obviously, for every set there is the property of being a member of that set. These properties will enable us to carve things up every which way we please. Nothing will ever be similar to

anything else, without at the same time being dissimilar to it in countless respects.

As far as the set-theoretic account of properties goes, the property of being a member of the set {computer, chair, Tilman} is on par with the property of being an Edmontonian for me. Clearly, however, we are looking for a certain economy and elegance in our descriptions of possible worlds. Not just any description will do. Lewis thinks this is a consequence of the way the worlds, and the entities constituting them, are put together. He believes that there is a certain privileged subset of the abundant properties - namely the sparse properties - which occupy an exalted position in terms of their fundamental importance.

> [T]hey carve at the joints, they are intrinsic, they are highly specific, the sets of their instances are ipso facto not entirely miscellaneous, there are only just enough of them to characterize things completely and without redundancy.[12]

Every possible world has its own set of sparse properties. In our world, for instance, it is physics which has begun the work of sorting out which of all the properties instantiated in our world are the fundamental ones, the properties which characterize our world completely and without redundancy. According to Lewis, who proposes to call these properties the natural properties, they are intrinsic to the entities possessing

---

12. Lewis, 1986, p.60

them. Lewis reminds us how much our daily philosophical practice depends on being able to make use of this distinction between natural and gruesome properties. In answer to those writers who have expressed scepticism concerning our ability to define the distinction between these two kinds of properties without presupposing it, he claims that the very necessity to constantly presuppose it is a strong reason to accept the distinction. And he goes on to accept is as primitive.

1.3 Propositions

Lewis takes propositions to be a certain kind of property. To be precise, a proposition is a property which is and can be instantiated only by entire possible worlds. Thus, as a common property is identified with the set of its instances, a proposition is identified with a set of possible worlds. Lewis's treatment of properties, most importantly the special case of propositions, will be of crucial significance in our critical discussion of his theory beginning in section 3. It is therefore imperative that we understand him correctly on this subject. To quote him:

> The proposition is the same thing as the property of being a world where that proposition holds; and that is the same thing as the set of worlds where that proposition holds.[13]

---

13. Lewis, 1986, pp.53-54

This, however, is not the case without qualification. Propositions which are sets of possible worlds do not discriminate between inhabitants of the same world. As Lewis points out, propositions, unlike sentences, are not supposed to be true or false relative to anything else but the world itself. The proposition that there exists someone whose name is Tilman is one such proposition. But the proposition that my name is Tilman fails under a definition of propositions as sets of worlds. If we want to capture the notion that propositions can, at least in some instances, be defined as the objects of thought, the kinds of entities which give us the content of what we know, then we will not be able to unpack this idea relative to possible worlds. The proposition that my name is Tilman is true for me, but false for a great many other people. What we need, says Lewis, are two different kinds of entities to fill the role of propositions: for the ones which are true or false relative to nothing but the world, or worlds, sets of worlds will do the job. But for the egocentric propositions we will need sets of possible individuals. Propositions, then, are either sets of possible worlds, or sets of possible individuals, as the case may be.

Section 2: The Ersatzist Challenge

As I have mentioned earlier, one of Lewis's most serious questions with regard to his theory concerns the possibility that its benefits may somehow be had without commitment to the sort of outsized ontology his particular approach requires. This is no idle concern. Indeed, there are a number of philosophers who would consider themselves, to a greater or lesser degree, modal realists without therefore feeling inclined to buy into Lewis's scheme of things. Saul Kripke, Rudolf Carnap, Richard Jeffrey, Alvin Plantinga, William Lycan, Robert Stalnaker, Peter van Inwagen, and Roderick Chisholm, to name just a few of the more prominent exponents, all share with Lewis the conviction that 'possible worlds' are in some way 'maximal objects', but do not endorse his view that those worlds are therefore spatiotemporal objects. These philosophers are united in the belief that there is only one concrete world, namely the one we are part of, and that possible worlds are abstract or non-actual entities, ways the concrete world might have been. Apart from this shared understanding, however, there is considerable disagreement about the nature of the possible worlds themselves. Some see them as a kind of linguistic entity, akin to maximal description of a state of affairs, some are less explicit about their exact nature, preferring instead to call them 'states of affairs', or 'propositions maximal with respect to entailment'.

In **OPW**, Lewis deals in detail with the cogency of the claims advanced by what he calls the 'ersatzers'. Since his treatment of the issue, more specifically the divisions of these claims into 'linguistic', 'pictorial', and 'magical' claims, has been followed by other writers in their discussion of the controversy, I shall follow him here.

2.1 Linguistic Ersatzism

Let us look at the linguistic ersatzer first. As we have mentioned, he is inclined to characterize possible worlds as sets of sentences, much in the way a novel might be. For the linguistic ersatzer, then, there is only one actual world, namely this one we happen to live in. Likewise, there is only one description of this world which is the correct one. This privileged description represents the 'way' this world is. All descriptions of ways worlds or things are which differ in any way from the actual one are incorrect descriptions of the one actual world. If they are maximal consistent sets of sentences, such descriptions constitute possible worlds. They are 'ways things might have been', complete descriptions of non-actual states of affairs.

On the face of it, this sounds like a useful way to approach possible worlds. We are dealing only with the kind of entities anybody would be able to be comfortable with. No exotic ontologies here. As far as the worlds are concerned, linguistic ersatzism asks no more than that we confine

ourselves to the one world we already know to be our own. Apart from that, there are only the constructs of the language we use to make our possible worlds, and this language can be almost any system of structures, as long as we can parse and interpret it. We might, for instance, let real numbers serve as words in our language, or spacetime points, as long as we stipulate their meaning. Lewis would even grant the ersatzer the use of the 'Lagadonian method'. In such a language, everything we might find in a given world just names itself. As Lewis points out, if we suppose that the concrete world consists of somehow 'elementary' particulars and universals, then we can let them serve as the vocabulary of our language. Consequently, pairs of particulars and universals would provide us with the atomic sentences of our worldmaking language, yielding a Lagadonian state description.

But all is not well in the linguistic ersatzer's paradise, and David Lewis is quick to point out three distinct problems with the approach which are as stubborn as air bubbles under a transparent plastic overlay. We can shift them around quite a bit, but it seems quite impossible to make them go away.

2.1.1 Primitive Modality

The first of these problems is the problem with primitive modality. Unless the linguistic ersatzer is content to offer nothing but a practical method of

exhibiting modal phenomena, if he has any desire to claim for his approach some explanatory function, he better not start assuming the very facts he is trying to explicate. Thus, if it is the nature of modality he wants to provide us an analysis of, modality better not be a part of the explanans. But this seems to be precisely what the linguistic ersatzer is forced to do in order to make his account work. We mentioned before that a possible world, in terms of a linguistic construct, has to be a maximal consistent set of sentences. It has to be consistent, because otherwise the description would portray an impossible world, not a possible one. And it has to be maximal, because otherwise it would not describe a complete possible world, but only a part of a possible world. A part, furthermore, which might turn out to be shared by two or more different possible worlds, once the missing sentences consistent with the incomplete set have been added.

In constructing his maximal consistent sets of sentences to serve as possible worlds, the linguistic ersatzer has to be sure there is not a single sentence which could be added to his description without making that description inconsistent. But how does he know that? Only by knowing that it is not possible to add a single further sentence. A maximal consistent description of a possible world is the maximal set of sentences which could possibly be true together. The linguistic ersatzer can offer us no

explanation of how he knows what that set is. Modality has slipped into his account as a primitive notion.

It is important in this respect to recognize that the ersatzer's problem does not arise because of the logical definition of maximal consistency for sets of sentences. In fact, Lewis's argument gains its force from this definition. It is because a set of sentences is inconsistent if and only if a contradiction can be derived from them, and maximal consistent if the addition of a sentence not already part of the set generates a contradiction that the ersatzer needs primitive modality to generate his worlds. The reason is that sets of uninterpreted sentences, however maximal consistent they might be, do not constitute possible ersatz worlds. It is their conversion into interpreted sets of sentences which is problematic. The ersatzer need not be disturbed by Lewis's claim about the 'reality' of possible worlds, but he will have to explain under which interpretations his sets of sentences come out as maximal consistent, and why. Given, for example, the two properties of positive and negative charge, it is possible to generate a maximal consistent set of sentences in the narrow logical sense of maximal consistency, in which a given electron has both positive and negative charge. But what if such a state of affairs is impossible as a matter of empirical fact? To get the modal facts right, we have to know what constitutes a possible interpretation for a given set of sentences. But

how does the linguistic ersatzer know that without taking modality as primitive?

2.1.2 Descriptive Power

The second problem has to do with the descriptive power of any language used for the making of possible worlds. Suppose, says Lewis, that we lived in a world in which, true to the beliefs of some ancient Greeks, the smallest constituent of physical reality was the atom. In our world, then, there are no such things as protons, electrons, or neutrons, to say nothing of quarks or neutrinos. It seems uncontroversially possible to imagine such a state of affairs. But how would we, in such a world, go about talking about a possible world such as the actual world happens to be? Even if we were not bothered by the fact that we have no names available to refer to the protons and electrons of that world, the fact that their properties are not instantiated in our world should give us descriptive problems. From the point of view of the atomistic world, our world possesses a number of properties which are neither instantiated in their own, simpler world, nor are they analyzable in terms of any other properties or relations in that world. Such properties Lewis calls 'alien'. They cannot be described in the atomistic world, because there are no words in the language of that world to describe them with.

Considering how easy it is to find an example like the one above, it would seem strange to want to resist applying

it to the actual world. It certainly seems possible that there is a world which exhibits properties alien to ours. Or, if we do not want to import possible worlds talk into this discussion, it certainly seems reasonable to say that there might be any number of alien properties which we do not presently know about, even some which we in principle will never be able to discover. Whichever way we put it, the linguistic ersatzer seems to face the dilemma of having to specify properties he does not know about. If he allows the possible existence of alien properties, he will be hard put to construct maximal consistent sets of sentences specifying the worlds which contain such properties.

In his recent discussion of ersatzism versus realism about possible worlds, Peter van Inwagen seems convinced by this argument[1]. He considers linguistic ersatzism defective, because it seems so easy for possible worlds to exceed the resources of a given language to describe them. However, I don not think that we should give up so easily. Lewis himself admits the possibility of 'Ramsifying' our linguistic ersatz worlds. A Ramsified ersatz world would be one which contained existential quantifications over properties such that these quantificational statements would be made true by there being certain alien properties. Our maximal consistent sets of sentences would then contain some which state or imply that there are, in fact, alien

1. Peter van Inwagen, 1986, **Two Concepts of Possible Worlds**, Midwest Studies in Philosophy XI, Minneapolis, University of Minnesota Press, p.201,

properties, and that they obey such-and-such nomological rules.

Lewis argues that this is not enough. He insists that we not only have to specify the roles of properties which are occupied, but also which properties occupy which roles. According to him, we can start with the description of a certain world, say ours, including the properties and their nomological roles, and then recombine the separate elements of that world in such a way that the resulting possible world will exhibit different nomological roles for the same properties. An example might help us to understand why this seems plausible to him. Suppose positive and negative charge turned out not to be intrinsic properties at all. Suppose they were, after all, not different determinates of one and the same determinable, but rather an extrinsic relation between particles. In his opinion, we would then have a situation where the same two properties of positive and negative charge had turned out to be obeying different nomological rules than we had at first believed.

Does this make sense? A further example will prove helpful here. In his discussion of the ersatzer's possible worlds, lewis claims that a given set of sentences may fall into inconsistency if the stated facts given a certain level of description imply a supervenient but inconsistent fact. He says:

We describe the spatiotemporal arrangement and fundamental properties of point particles in infinite detail - and lo, we have implied that there is a talking donkey, or we have fallen into inconsistency if we also say explicitly that there is not.[2]

If we should accept this as plausible, as maybe we should, why would we not also accept that a description of the spatiotemporal arrangement of point particles and some fundamental properties will imply the existence of other properties? It seems that such properties are defined precisely by their distribution and nomological role. We postulate their existence in order to account for the interesting regularities in the behaviour of elementary particles not so far subsumed under known properties. Similarly, were we to give a local description of point particles and their fundamental properties excepting therefrom the property of positive or negative charge, such a description would imply the property of positive or negative charge if the description is of particles in our world. To say that positive and negative charge might not be the intrinsic properties we take them to be, but might instead be extrinsic, relational properties, is to imply that they will share with the former properties nothing but the name.

---

2. Lewis, 1986, p.155

Consider a world where charge, as Lewis suggests, is indeed an extrinsic property, whatever that might be taken to mean. It seems at least acceptable that such a property would be relational. But then it would be possible to give a consistent description of the states of affairs in this world, omitting from the description only the name of the particular relational properties. It seems unclear how anything the sentences of that description might jointly imply would cause us to conclude the existence in that world of the properties of positive and negative charge as we know them, or alternatively, as Lewis would have us believe, of those same two properties, albeit obeying different nomological rules than in our world.

Lewis is clearly led into this confusion by relying too much on the productivity of what he calls the 'principle of recombination'. Not only does this principle allow him to patch together 'things' from different worlds to create new and different worlds, he believes that he can also patch together properties divorced from their nomological roles. As he states: "We can distinguish our world from one in which, say, one of the quark colours has traded places with one of the flavours."[3]

But can we? Recall that, for Lewis, a property is the set of all its possible instances. Thus, the property of being a flying pig is the set of all flying pigs in all

3. Lewis, 1986, p.162

possible worlds. Since flying pigs are implied by complete low level descriptions of the distribution of point particles and the relations holding between them , as Lewis states in his discussion regarding worldmaking languages, we can also say that the property of being a flying pig is the set of all complete descriptions of point particles and the relations holding between them in spatiotemporal regions coextensive with flying pigs. These descriptions will not mention flying pigs explicitly, but the nomological rules governing the distribution and interrelation of the point particles will imply that there are flying pigs.

But if the principle of recombination allows the divorce of properties from their nomological roles, then it will also allow the divorce of lower level descriptions from their higher level implications, simply because the property of being a flying pig can then no longer be said to follow from a given lower level description with any necessity. After all, the description might be of a world where the property of being a flying pig was divorced from the nomological role played by it in other worlds containing flying pigs. Instead, in this case, the same nomological role was instead played by the property of being a flying mouse.

It seems that if Lewis wants to retain the power of his principle of recombination, he will have to soften his stand

on implication, thus weakening his criticism of the descriptive power of linguistic ersatzism.

2.1.3 Indiscernible Individuals

This is Lewis's third objection to linguistic ersatzism. It concerns indiscernible entities which, according to linguistic ersatzism, are just multiple actualizations of one and the same description. But it seems possible to speak of a universe consisting of nothing but a perfect crystalline lattice infinite in all directions, or of a universe of eternal recurrence. The point here is that if we can sensibly speak of a possibility containing a number of different individuals, as with the lattice or the individuals populating the universe of eternal recurrence, then our description should somehow be able to refer to the different possible individuals without conflating them all into one. Linguistic ersatzism, it seems, cannot avoid such conflations. If there are two indiscernible possible individuals, there will only be one description, thus misrepresenting the facts of modality.

2.1.4 Summary

To recapitulate: linguistic ersatzers face three major difficulties with their position, one of which they might succeed in sidestepping, but two of which will remain recalcitrant. These problems are:

1. Linguistic ersatz possible worlds lack the resources to correctly mirror all the modal facts there are. This

> refers to the fact that there might be possible worlds so rich as to make it impossible for any language to give explicit descriptions of them, especially in the case of alien properties. This is the problem which might go away, provided we are prepared to consider the nomological roles of properties essential to them, and to give 'Ramsified' descriptions of our possible worlds.
>
> 2. Linguistic ersatzers need primitive modality to make sense of their enterprise. For them, possible worlds are maximal consistent sets of sentences. But a maximal consistent set of sentences is just the maximal set of those sentences which could all be true together. This problem will not go away.
>
> 3. The descriptions which constitute the possible worlds of linguistic ersatzers will not be able to account for indiscernible possible individuals, as there will always be only one description remaining ambiguous between the many possible indiscernible individuals. This is the other problem that will not go away.

David Lewis is correct in criticizing linguistic ersatzism for these reasons, quite apart from any operational benefits its advocates might claim for it. What remains to be seen is whether there is in fact another theory of modality which would offer us all the advantages of the linguistic ersatzist approach without at least some of the drawbacks we just discussed. Let us turn, therefore, to the next ersatzist enterprise Lewis mentions in his discussion.

2.2 Pictorial Ersatzism

As Peter van Inwagen has pointed out[4], pictorial ersatzism is a theory no one in fact holds. It seems therefore hard to understand at first why Lewis would find it worthwhile to devote a portion of his book to a detailed discussion of this view. The reason for its inclusion, I think, becomes clear once we have acquainted ourselves with his perspective. Pictorial ersatzism turns out to be a kind of primitive version of his own brand of modal realism, and he needs to prove its shortcomings in order to justify the further step which transforms the pictorial model into the full fledged realism of his theory. In the following, I will therefore follow his treatment of pictorial ersatzism and his definition of pictorial representation.

In essence, pictorial ersatzism is an answer to the problems of linguistic ersatzism. Given that the linguistic model ran into difficulties with its attempt to represent possible worlds by way of maximal consistent descriptions, maybe we ought to look for another way of construing the representation relationship. Linguistic ersatzism tries, and fails, to represent as a story would. It seems as if pictorial ersatzism could maybe avoid this problem-riddled approach. According to this alternative, possible worlds are pictures. They represent the actual world the way a very elaborate and detailed picture would. The relationship here

4. van Inwagen, 1986, p.201

is one of isomorphism. Of course, it would be false to imagine this to be the kind of isomorphism we are acquainted with from the way pictorial representation works in the world we live in. Pictures, models, maps, etc., as they are ordinarily used require extensive conventional interpretation rules to bridge the substantial gaps between the representation and the actuality to which they are supposed to be isomorphic. The necessity of such interpretations makes conventional pictures hybrid forms between linguistic and pictorial representation. Lewis suggests that in order to discuss an approach to possible worlds which would avoid the particular shortcomings of linguistic ersatzism, we might as well choose something as different from it as possible. Thus, pictorial ersatzism, as he conceives of it, represents entirely by isomorphism with respect to the world it represents. In the case of our own world, the pictorial ersatz world represents the actual one by consisting of parts, properties, and relations which stand in a precise one to one relationship to the respective parts, properties, and relations in our world. The pictorial ersatz world is isomorphic to our world with respect to every detail. As Lewis says:

> We want no projection into too few dimensions, no omission of hidden or minute parts, no spurious structure that is in the picture but plays no part in the representing.[5]

5. Lewis, 1986, p.166

Obviously, such a representation is an idealized version of whatever we might be acquainted with in the way of pictorial representation. But we might as well stipulate that there be not a single part of the world which is not mirrored in the pictorial ersatz world in exactly the corresponding position, relation, dimension, etc., to make it a complete representation of the original. The short answer to the question: 'isomorphism with respect to what?' is for purposes of our discussion: 'with respect to everything.'

2.2.1 Primitive Modality Revisited

Pictorial ersatzism does not have the worries about consistency which beset linguistic ersatzism. Given our 'idealized' definition of 'isomorphism with respect to everything', there can be no inconsistency in the representation. Nevertheless, isomorphic representation seems to need to import the resources of primitive modality in order to make sense. Lewis reasons as follows: the kinds of ersatzist programs examined here all agree that there is one and only one actual concrete world, namely the one we live in. Likewise, there is only one correct, pictorial representation of this world (even if we might allow the possibility of multiple indiscernible correct representations). This representation, being truly isomorphic to the actual world, represents it accurately. All other pictorial representations, since there are no

worlds for them to be isomorphic to, fail to represent anything. What then, Lewis asks, might it mean for a pictorial representation to represent a talking donkey? There are no talking donkeys in the one concrete world for any representations to be isomorphic to. So what makes the pictorial representation a representation of a talking donkey? Lewis says:

> ...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, and it couldn't have been isomorphic to any part of the concrete world that was not a talking donkey.[6]

But what if we claimed that ersatz entities are defined, not by their isomorphism to concrete entities but by their structural identity with sums of ersatz parts which in turn are isomorphic to parts of the concrete world? Lewis counters that this would require us to know the structural definitions of such things as talking donkeys and philosophizing cats before we could know which combination of ersatz parts isomorphic to parts of the concrete world would be a representation of one of those possible entities.

2.2.2 Conflation of Indiscernibles

The second problem which pictorial ersatzism shares with the linguistic rival is its inability to distinguish indiscernible individuals. Just as any description which is true of one actualized individual is equally true of any

---

6. Lewis, 1986, p.168

other individual indiscernible from the actualized one, so any picture which is isomorphic to an actualized individual 'with respect to everything' will be equally isomorphic to any other individual indiscernible from the first. Also, given that we have a number of indiscernible pictorial representations, one single individual which actualizes one of them by being isomorphic to it will necessarily actualize them all.

2.2.3 The Collapse of Pictorial Ersatzism

Of course, the real reason why there is nobody advocating pictorial ersatzism is not to be found in the objections we have so far considered. Rather, there are no pictorial ersatzers simply because a pictorial ersatzer is actually a modal realist who has not yet become conscious of that fact. In a way this was implicit in the description of the theory all along. For if pictorial ersatzism represents by isomorphism, and if that has to be considered a very special, very idealized kind of isomorphism in order to be able to represent an existing world fully and correctly, we are led to the conclusion that a pictorial model represents by being exactly like its actualized original.

The point here is the same one we mentioned before: representation by isomorphism which is not based on an isomorphism with respect to all methods of investigating will need some interpretive rules to specify the ways in which the isomorphism is to be understood. Lewis correctly

objects to anything less than isomorphism 'with respect to everything' (his idealized pictorial representation) on the grounds that it would mix elements of linguistic ersatzism with pictorial ersatzism, thus importing the problems of linguistic ersatzism into the very theory devised to sidestep them.

By definition, then, the isomorphism is defective if anything that is present in the actual world is not there in the model to represent it. But then it becomes difficult to imagine how we might conceive of the difference between a pictorial representation and the actual world it represents. Or, to come to the point Lewis was driving at all along: once our pictorial ersatz worlds become detailed and complete enough to satisfy the requirements for isomorphic representation set out above, they have become possible worlds in the modal realist sense. To try and deny this is to try and insist in some magical difference between them, something (but not existence, not concreteness, or actuality) which separates the actual world from its truly isomorphic representation. After all, the pictorial ersatzer wants to agree when we demand that nothing must be left out by the representation if it is to count as truly isomorphic. Unfortunately, if it does leave something out, then it does leave something out, so the purported difference can't be a difference, or else the theory collapses. But if there is no difference, the theory collapses as well. Pictorial

ersatzism, we see, is incoherent and collapses into modal realism by virtue of its own premisses.

2.3 Magical Ersatzism

In his discussion of magical ersatzism, David Lewis takes aim at the great majority of modal logicians who do not consider themselves purely in the linguistic camp. Although he remains quite noncommittal, electing to discuss entities like 'abstract simples', and relations like 'selection', it is not difficult to discern the thrust of his attack. We need only replace the 'abstract simples' with 'propositions', and the 'selecting relation' with the 'truth value selection relation', to understand who the targets of Lewis's criticism are meant to be: anyone who chooses to talk about modality in terms of 'states of affairs', 'ways things might be', 'possibilities', 'propositions', or 'structures', without wanting to specify exactly what those entities might be composed of, save that they are not spatiotemporally extended or set theoretical constructs on spatiotemporally extended objects. What they do want to say is that to be true or to be the case is a relational property of those entities. There is exactly one concrete world, namely ours, and it is this world which assigns a particular proposition its truth value, or makes a particular state of affairs be the case. Likewise, the concrete world makes exactly one of all of the possible worlds (seen as propositions closed under entailment) actual. It can do so because propositions qua propositions ar neither true nor false. It is the concrete, actual world which selects the truth value for each proposition. The

proposition that flying pigs exist, which in this form would come out true on Lewis's construal of possibility and propositions (there is a world where pigs fly), is decidedly false according to this view. The one and only actual world simply does not contain flying pigs. However, had our world been differently arranged, that proposition might have been selected true. All this has nothing to do with the structure of our language, our mind, or even with our existence. Propositions are necessarily what they are. It is their relation to the objectively given actual world which is contingent.

Lewis proposes to be as accommodating as possible. He does not want to take issue with the somewhat vague definition of what propositions are supposed to be. He even agrees to allow that the 'truth value selection' relation be taken as primitive. But, says he, we ought at least be able to give an account of how it is we perceive its effects on a given proposition. Maybe it "would be unfair asking the ersatzer to define his primitive. But I am not demanding definition, only classification."[7] And he goes on to illustrate this with a point from his own theory: he might take predicates of distance as primitive, but he is emphatic about the fact that distance is an external relation, not affecting the intrinsic properties of the relata in any way, nor depending on them. If the magical ersatzer is to avoid an account of modality without any explanatory power, he

7. Lewis, 1986, p.176

should at least be able to tell us if the mysterious relation which the actualized world bears to propositions is an external one or an internal one.

2.3.1 External or Internal?

Suppose the ersatzer claimed that the relation in question between the actual world and a given proposition was external. That would mean that it would hold independently of any intrinsic properties the world or the proposition might have. Presumably this would suit the magical ersatzer quite well, since it would make his reluctance to explain the nature of his elements more palatable to us. Given that the relation by which the concrete world selects a proposition is an external one, all we need to worry about is the position of our propositions in our relational system, since this is the only determining factor in the truth value selection relation. The problem here is, most obviously, that the ersatzer cannot want this relation to be independent of what goes on in the concrete world, since the concrete world was posited as grounding the very relation in question. So he has to say that there is a necessary connection between a certain particular state of the actual world and a particular proposition. But this is not due to any structural or whatever other similarity between the world and the proposition. Nothing in the nature of the proposition determines the relationship. It is an external relation. Now, if it were truly an external

relation, totally independent of any intrinsic properties the proposition might have, we should be able to understand the relation even if we construed all propositions to have identical intrinsic proprieties. On this view, the proposition that there are flying pigs, and the proposition that there is a cat on a mat, might as well be treated as if they were intrinsically identical. Consequently, it becomes possible to suppose that the concrete world could have been intrinsically exactly the way it actually was when it bore the particular truth value selection relation to the proposition that there is a cat on a mat which made this proposition true, and yet it could have made the proposition that there are flying pigs true instead.

No ersatzer can admit to such a state of affairs and at the same time maintain that there is one and only one actual concrete world which makes certain propositions true and others false due to the way that world is intrinsically structured.

There seems to be only one way out of this dilemma, and that is to understand the relevant relation as a primitively modal one. But what have we managed to explain about modality if we need to assert that there is a necessary relation between the actual, concrete world and the proposition made true by it? Lewis is right in claiming that this amounts to postulating a magical relation.

So what if we were to assume the relation to be an internal one? In other words, the 'truth value selection relation' might be dependent on the intrinsic properties of the world and of the propositions. As far as the actual world is concerned, we might be able to make sense of that. But what about the propositions? To talk about their intrinsic properties cannot mean that they have spatiotemporal, or mereological, or even set-theoretical structure. This much the ersatzer has already committed himself to. But then, what might it mean to speak of intrinsic properties? And how would we know about them, given that propositions have no spatiotemporal extension or location, and are therefore not causally related to us? If there is nothing that I can say about the intrinsic properties of a proposition, simply because there is no causal connection between me and the proposition, no telescope with which to view this elusive entity, if all I can say about the proposition is that it stands in the 'truth value selection relation' to such-and-such an object in the actual world, what sense does it make to call that relation internal?

Let us imagine some quasi Platonic world in which objects bear a very special relationship to their counterparts in the realm of the forms. We claim that this relationship is internal. Certain objects bear it to certain forms by virtue of their respective intrinsic nature. However, we have no direct access to the realm of the forms.

We cannot simply behold the forms and distinguish them one from the other. Our only way of knowing them is by virtue of the special relationship they bear to certain objects here on earth. Let us call this relationship 'participation'. So when we want to refer to any form in particular, we have to individuate by saying: the form that such-and-such a set of object participate in.

The question Lewis asks the ersatzer is: do you think that this means you understand the 'participation' relation? It seems that to understand a certain relation would entail that we know how to point it out when we encounter an instance of it. But how could we ever encounter an instance of the 'participation' relation, taking it to be a dyadic relation for the sake of simplicity, if one of its relata can only be distinguished by means of using the notion of participation in the first place? A truly internal relation would allow us to recognize it by virtue of certain of the intrinsic properties of the relata. The property of bearing the 'participation' relation to a certain set of objects cannot itself be considered an intrinsic property of the forms, if we want to avoid circularity. However, we know of no other individuating features of the forms. To understand the relation of 'participation' without knowing the intrinsic properties of the things it is borne to must therefore be a magical ability. Lewis is right in dismissing magical ersatzism as flawed.

2.4 Summary

The arguments Lewis advances in his attack on the ersatzist enterprise seem, for the most part, to be sound. In some cases the ersatzer might be able to salvage his position by some crafty shifting of ground, but the substantive issues cannot be argued away. In particular, there are three specific problems which dominate the issue:

1. The ersatzer, particularly the linguistic ersatzer, encounters difficulties concerning reference. This becomes apparent for instance with regard to the naming of alien properties, etc..

2. The ersatzer is hard pressed to make his account of the relationship between the actualized world and the representing entities sound any less than magical, unless he opts for the linguistic approach, which has other drawbacks.

3. The ersatzer seems unable to get rid of recalcitrant primitive modality of the most obnoxious kind, be it hidden in the notion of maximal consistency with respect to worlds, as with the linguistic ersatzer, or show itself as a necessary connection between the world and propositions, as in the case of the magical ersatzer.

Lewis gains much mileage by concentrating his criticism of the ersatzist program on these three recurring problems. It seems only fair, therefore, to examine his own position with regard to these issues. Is his theory of modal realism safe from the objections he raises against the ersatzers? If so, what are the features of his theory which are responsible for this? If not, how vulnerable is his position, and why? In the following, we will turn to an examination of this interesting set of questions.

Section 3: The Magic of David Lewis

As we have seen in the last chapter, most of the arguments David Lewis employs in his critique of modal theories competing with his own gain their strength from the inability of their proponents to explain the nature of the relations they postulate as holding between the actual world and the representing class of entities. In fact, they cannot even classify these relations to tell us if they are external or internal. In addition, all of them do, in some form or other, presuppose modality as primitive in their respective theories, a fact which makes them of little value for any analysis of the notion of modality.

Lewis does not regard any of the arguments against the ersatzist enterprise as conclusive proof that one ought not to engage in it. What he does want to claim is that his own modal realism combines any of the practical advantages of the other theories in terms of the work they can do with a much needed analysis of the nature of modality itself, thus for the first time enabling us to know what it is we are actually talking about when we use modal terms. This is what he considers to be the great advantage of modal realism, and this is what makes modal realism worth the ontological price in his estimation.

Needless to say, for this claim to retain its force, he needs to be extremely miserly in his use of primitive terms,

and most of all, he has to beware of the deadly trap of primitive modality. I believe he fails in both respects. Furthermore, I believe that he really did not have a chance from the start. The problems he so aptly exposes in the linguistic, the pictorial, and the magical ersatzer do not arise because those theorists fail to subscribe to modal realism. Nor are they fundamentally different problems from case to case. As I will try to demonstrate in the following pages, the ersatzer's difficulties with modality are rooted in the fact that they are all metaphysical hyper-realists, if not modal hyper-realists in the correspondence theory of truth sense of the word. Mutatis mutandis, David Lewis does not avoid their difficulties, he merely imports them into a vastly expanded ontology.

Let us begin with what might be called the 'root of the problem'. It is not new. Hilary Putnam has tried to bring it to our attention since 1977.[1]

3.1 The Indeterminacy of Reference

In the dozen years which have passed since Hilary Putnam first presented his disquieting conclusions to the American Philosophical Association, and to the Association for Symbolic Logic, a number of people have made attempts to shield model-theoretic semantics from the impact of his conclusions. As several of them have tried to do so without

---

1. Hilary Putnam, 1977, **Realism and Reason**, Proceedings of the American Philosophical Association, 50, pp.483-498,

clearly understanding the force of his objection as well as its extent, it will be useful to set out as clearly as possible what 'Putnam's Theorem' proves and why.

3.1.1 The Model-Theoretic Component

Here, first of all, the predicament in Putnam's own words:

> The predicament only is a predicament because we did two things: first, we gave an account of understanding the language in terms of programs and procedures for using the language (what else?); then, secondly, we asked what the possible "models" for the language were, thinking of the models as existing "out there" independent of any description. At this point, something really weird had already happened, had we stopped to notice. On any view, the understanding of the language must determine the reference of the terms, or, rather, must determine the reference given the context of use. If the use, even in a fixed context, does not determine the reference, then use is not understanding. The language, on the perspective we talked ourselves into, has a full program of use; but it still lacks an interpretation.[2]

According to Putnam, then, the problem which he wants us to understand is one which is implicit in the received and accepted view of model-theoretic semantics. The problem is the way we are trying to see language and its relation to the world. As long as we hold on to a theory of meaning which sees the syntactical features of a language as non-contributive to semantics, we will have his problem. But why, exactly?

---

2. Hilary Putnam, 1980, **Models and Reality**, The Journal of Symbolic Logic, Vol.45, No.3, p.481

The simple answer is that the pairing of uninterpreted strings of symbols with uninterpreted model structures does not yield meaning. What seems to have gone unnoticed in much of the work in semantics before Putnam dropped his bomb was the fact that the models used to provide the satisfaction conditions for sentences are themselves nothing but meaningless structures. Without being given an interpretations they are just neutral structures composed of sets of abstract entities. In mathematics, this fact and a range of problems arising from it had already been known since the publication of the Lowenheim-Skolem theorem:

> If a countable collection of sentences in a first-order formalized language has a model, it has a denumerable model.

Before then, mathematicians might have considered it possible to come up, for instance, with an axiom system for the real numbers. Such a system would have to distinguish the real numbers from the integers and do so with a countable number of terms (presupposing that if humans could create such a system it would have to be limited to countably many axioms). But recognizing that the set of axioms, as we said before, would start out as a set-theoretic structure without meaning, a purely syntactic construct consisting of countably many terms, Lowenheim and Skolem realized that it would always be possible to supply a set of countably many objects to satisfy all of the axioms. This pairing could be done prior to any assignment of meaning. In other words, any countable set of axioms

purporting to state all and only the truths about the real numbers could always be assigned an alternative interpretation according to which it stated all and only the truths about the integers. Nothing about the structure of the axiom system would fix its reference exclusively to the real numbers.

However, Putnam is not concerned with applying the Lowenheim-Skolem theorem to the semantic theory of natural languages. He simply noticed that the general insight underlying the Lowenheim-Skolem theorem, namely the fact that truth in a model can be preserved while the interpretation is radically changed, has important ramifications for model-theoretic semantics as well.

3.1.2. The Impact on Model-Theoretic Semantics

Putnam's strategy is as simple as it is effective, based on the standard model-theoretic definition of meaning:

> 1. The meaning of a sentence is a function which assigns a truth value to that sentence in each model (possible world).
>
> 2. The meaning of a term (predicate) is a function which assigns a referent to that term (predicate) in each model (possible world).

Putnam's argument turns on the realization that the peculiar model-theoretic property which gives rise to the Lowenheim-Skolem theorem, makes the above mentioned definitions of model-theoretic semantics insufficient for a consistent theory of meaning. This is because truth

underdetermines reference in model theory. One can change the reference of the parts of a model while preserving the overall 'satisfaction in the model' of a given set of sentences. Putnam shows how this can be done in the case of natural language, by proving that one can preserve the meaning, i.e. the truth value assignments, of a sentence like (1) 'A cat is on a mat' for all possible worlds, and at the same time change the meaning, i.e. the referents, of 'cat' and 'mat' so radically that 'cat' now refers to cherries, and 'mat' to trees. As he says:

> I shall show that sentence (1) can be reinterpreted so that in the actual world 'cat' refers to cherries and 'mat' refers to trees without affecting the truth-value of sentence (1) in any possible world.[3]

There seems to be little use in trying to paraphrase Putnam's ensuing demonstration, but at the same time we need to look at it in detail in order to understand the attempts that have been made to blunt his attack and the reasons why they fail. Consequently, I shall quote his own exposition here, at the risk of boring some readers. He continues:

> The idea is that sentence (1) will receive a new interpretation in which what it will come to mean is:
>
> (a) A cat* is on a mat*.
>
> The definition of the property of being a cat* (respectively, a mat*) is given by cases, the three cases being:
>
> (a) Some cat is on some mat, and some cherry is on some tree.

---

3. Hilary Putnam, 1981, **Reason, Truth and History**, Cambridge, Cambridge University Press, pp.33-34

(b) Some cat is on some mat, and no cherry is on any tree.

(c) Neither of the foregoing.

Here is the definition of the two properties:

Definition of 'cat*': x is a cat* if and only if case (a) holds and x is a cherry; or case (b) holds and x is a cat; or case (c) holds and x is a cherry.

Definition of 'mat*': x is a mat* of and only if case (a) holds and x is a tree; or case (b) holds and x is a mat; or case (c) holds and x is a quark.

Now, in possible worlds falling under case (a), 'A cat is on a mat' is true, and 'A cat* is on a mat*' is also true (because a cherry is on a tree, and all cherries are cats* and all trees are mats* in worlds of this kind). Since in the actual world some cherry is on some tree, the actual world is a world of this kind, and in the actual world 'cat*' refers to cherries and 'mat*' refers to trees.

In possible worlds falling under case (b), 'A cat is on a mat' is true, and 'A cat* is on a mat*' is also true (because in worlds falling under case (b), 'cat' and 'cat*' are coextensive terms and so are 'mat' and 'mat*'). (Note that although cats are cats* in some worlds - the ones falling under case (b) - they are not cats* in the actual world.)

In possible worlds falling under case (c), 'A cat is on a mat' is false, and 'A cat* is on a mat*' is also false (because a cherry can't be on a quark).

Summarizing, we see that in every possible world a cat is on a mat if and only if a cat* is on a mat*.

And in the appendix of Reason, Truth and History Putnam adds a detailed proof of 'Putnam's 'Theorem':

Let L be a language with predicates F, F,...F (not necessarily monadic). Let I be an interpretation, in the sense of an assignment of an intension to every predicate of L. Then if I is non-trivial in the sense that at least one predicate has on extension which is neither empty or universal in at least one possible world, there exists a second interpretation J which disagrees with I, but which makes the same sentences true in every possible world as I does.

According to the proof Putnam supplies, which I will not bother to replicate here, it is a fact of model theory that it cannot be used to provide a consistent theory of meaning. This does not mean that model theory cannot play some part in a successful theory of meaning. It means that as long as we entertain the view of language as an initially meaningless syntactic structure in need of an 'interpretation', and expect 'truth' to be definable by a function which specifies 'satisfaction conditions' for sentences of a language in terms of the set-theoretical structure of a 'model' of entities and their relations, we will underdetermine the reference conditions of the terms of the language. But to qualify as serviceable, any theory of meaning must at least specify some non-trivial relation between the meanings of the parts of sentences and the meaning of the whole such that changing the meanings of the parts will also change the meaning of the whole. As we have seen, model-theoretic semantics, with its definition of meaning of sentences in terms of truth value, cannot help us there.

For the same reasons Putnam's result cannot be avoided by simply specifying all the 'correct' referents of a language. At least not as long as we maintain that the meaning of a sentence is the function which assigns a truth value to the sentence in every possible world. Putnam has

already demonstrated that we can assign a totally preposterous interpretation to the terms of a language and still preserve all the truth values of a prior, more 'natural' interpretation. If, therefore, an 'incorrect' interpretation comes out as being as true as a 'correct' one, we have no grounds for rejecting the former and accepting the latter. Truth in the model was, after all our criterion. And it is satisfied equally well by both.

Obviously, this is not a conclusion any self-respecting metaphysical realist will be prepared to live with. Consequently, a number of attempts have been made to discredit Putnam's line of argument. For the purposes of this thesis, I will concentrate on the ones which seem pertinent to our discussion of possible worlds.

3.1.3. Meaning Just Ain't In The Head

David Lewis begins his counterattack against the Putnam theorem as follows:

> Hilary Putnam has devised a bomb that threatens to devastate the realist philosophy we know and love. He explains how he has learned to stop worrying and love the bomb. He welcomes the new order that it would bring (RT&H, Preface). But we who still live in the target area co not agree. The bomb must be banned.[4]

How does he propose to accomplish this? For Lewis the way to the solution seems obvious: "realism needs realism" is the motto which he proposes to inscribe on the shields of

4. David Lewis, 1984, **Putnam's Paradox**, Australasian Journal of Philosophy, Vol.62, No.3, p.221

the defenders of conventional model-theoretic semantics. In other words, anyone who would seriously want to advocate realism should also claim that the 'objective world', which he believes to be out there independent of his mind, acts to restrain rampant indeterminacy of the Putnamian kind by possessing a kind of 'social' organization simpliciter. This inherent organization divides 'reality' into things and classes according to a scale of 'eligibility' based on 'objective sameness and difference in nature'. Thus, according to Lewis, it is the case that among all the many things and classes which make up our world, some elite few are of the aristocratic class, and most others are unimportant commoners. The aristocrats are carved at nature's objectively real joints, and wholly so. The rest are gerrymandered to a greater or lesser extent, thus providing us with a convenient scale. It is the task of physics to discover exactly what the nobility consists of. In the opinion of David Lewis, we are well on the way of accomplishing this. Mass, charge, quark colour and flavour are already certain to be of the coveted blue blood. Equally certain, cats, pencils and puddles can be accorded moderate elite status. And even strange properties like Nelson Goodman's 'grue' can still be considered higher on the scale than Putnam's 'cat*' or 'mat*'.

Given that such a state of affairs really is the case, claims Lewis, the discovery of what is the truth about the world becomes a non-trivial undertaking once again. Not all

interpretations will respect the objective samenesses and differences, the joints in nature etc., that there actually are. Thus, not all interpretations are on a par with respect to their truthvalue assignments, just as long as the sets of sentences they generate come out satisfied by the same model. Interpretations will be graded according to how successfully they account for the way the world is 'objectively'. And it becomes possible again to speak of correct interpretations and incorrect ones. At least for Lewis.

Granted, we would intuitively agree that such crazy semantic reassignments as Putnam's 'cat*' and 'mat*' should be ruled out by a successful semantic theory. I don't think that Putnam would demand any less. The question was and is how to accomplish this, and the claim to be disproved is that model-theoretic semantics is not up to this task. What then of Lewis' objective restraints? Do they offer the way out of the dilemma?

First of all, a theory which advocates a hierarchical ordering of the entities which constitute reality according to their more of less 'elite' status would strike any but the most dyed-in-the-wool metaphysical realists as ad hoc. What after all, except its utility in possibly saving model-theoretic semantics from falling into inconsistency, could be cited as further evidence of such an objective ordering? Lewis seems to think that levels of reductive description

should be taken as evidence. Thus, because we can give a reductive analysis of cats, pencils and puddles in terms of mass, charge, quark colour and flavour etc., that makes the latter more aristocratic ingredients of reality than the former. But why should we accept this? Even if we were to agree with him on the objective existence of a hierarchical ordering of all the entities of reality in terms of their 'naturalness', it seems entirely arbitrary which level to assign top status to. What if we granted Lewis his claim that nature is 'really' carved at her joints at the level of cats, and pencils, and puddles? Or maybe we should leave out the pencils, as Lewis should have, considering his criterion of 'discriminatory classifications not of our making'. So what if cats and puddles and cherries were the true aristocracy of existents? Couldn't we claim that the reductive analysis of these privileged entities in terms of more 'basic' elements obscures rather than outlines the objective joints in nature, analogous to the way in which the macroscopic world of our naive sensory acquaintance is lost in microscopic description of it which does not stipulate a priori which implication relations are supposed to hold between the two levels? The fact that physicists are currently quite interested in what Lewis calls the 'fundamental' properties of nature does nothing in and of itself to bestow any kind of privileged status on these entities, whatever they may be. Lewis does nothing to explain to us why the fact that physicists give priority to

the search for reductive analyses of natural phenomena ought to be taken as a reason to consider lower level entities more 'elite' than higher level entities.

Secondly, and more importantly, even if there was indeed such a well-ordering objectively given in nature, this would not rule out moderate indeterminacy, as Lewis himself admits several times in his critique of Putnam: "I take it that the existence of moderate indeterminacy is not to be denied."[5]

Lewis does not seem to much concerned with this fact. But why? As far as model-theoretic semantics is concerned, it would seem that any kind of indeterminacy would be enough to bring the enterprise to a halt. After all, the hope which attaches to a theory of semantics based of model theory is the hope of being able to treat semantics with rigour, of being able to state the truth conditions of sentences precisely, and to demonstrate the essential connection between the assignment of truth values to sentences and the assignment of referents to the terms of the language. Any theory of semantics which is prepared to let the referents of the terms of the language vary while holding fixed the truth value of the sentence they are part of, will need some further mechanism which ensures that the meaning of the parts cannot change without thereby changing the meaning of the whole, if it is to be an adequate theory of meaning.

---

5. Lewis, 1986, p.223

But, as we have said before, the model does not interpret itself. What interpretation it receives it receives from the assignment of its own parts as referents to parts of sentences. Anything which could be advanced as the desired mechanism would simply be an addition to the sentences of the language, and as such satisfiable in a multiplicity of different models. This is true whether we see the saving constraint as 'causal', (replaced by 'causal*'), 'natural' (replaced by 'natural*'), 'elite' (replaced by 'elite*'), or any other theoretical device we might try to suggest. And as Lakoff observes[6], it will not help to admit that, although theoretical devices will not help us to rescue our theory of meaning, there is in fact a constraint operation in the world itself which our theory merely conforms to. This is because designating some presumably 'objectively real' constraint OC as the true guarantor of correct reference would entitle us to claim that "OC is the real constraint" is true in our language. And again we would be able to assign any number of referents to the term OC which would make the sentence come out equally true.

Lewis does not seem to appreciate properly how much of the theoretic approach he loves has been destroyed by the 'bomb'. Putnam could cheerily grant him all his 'natural' properties as given and then proceed to reassign referents to the terms of the language by switching only among the

---

6. George Lakoff, 1987, **Women, Fire, and Dangerous Things**, Chicago, Chicago University Press, p.240

properties Lewis accepts. This would not result in the rampant indeterminacy of the 'cat on the mat' variety, even though one would want to know at which level indeterminacy should be called rampant instead of moderate. Lewis says:

> [O]verall eligibility of reference is a matter of degree, making total theory come out true is a matter of degree, the two desiderata trade off... The terms of the trade are vague; that will make for moderate indeterminacy of reference; but the sensible realist won't demand perfect determinacy...[7]

Given that such 'moderate' indeterminacy might still generate many thousands of examples it seems hard to avoid the conclusion that Lewis' amended theory of reference would lapse into inconsistency.

3.1.4. Does Lewis Have No Models?

It has been suggested that Lewis' modal realism is not vulnerable to any of the attacks on model-theoretic semantics, simply because he doesn't have any models in his theory. David Lewis, so this story goes, doesn't have models. He only has sets of sentences and real worlds. This is in fact precisely the advantage he enjoys over modalists of the other persuasions. His modal realism allows him to let the symbols of his language refer directly to the worlds themselves, and to the things they contain.

This sounds enticing and initially quite promising. Unfortunately, as George Lakoff observes[8], even if the above

---

7. Lewis, 1984, p.228
8. Lakoff, 1987, p.253

characterization of Lewis' semantic theory were correct, Putnam's bomb would damage it. And here is why:

Let us assume that Lewis has something we might call a 'direct' theory of meaning. According to it, nothing is relevant but the pairing of symbols with objectively existing things in the world. This process does not in any way depend on being perspicuous to any thinking being. The relation between the symbols and the things is an objective fact. For instance, we could follow Lewis and say that the term 'world' just means the mereological sum of all the things which are spatiotemporally related to each other. But to rank as a theory of meaning, our theory must not only assign referents to symbols, but also meanings to sentences and parts of sentences. Since this cannot be dependent on the presence of any symbol users, it follows that the semantic relationship between the individual terms and the whole sentences must be an objective one. This, says Lakoff, begins to sound a lot like model theory, and not by accident. The reason is that a theory which assumes an objective relationship between existents as part of a world and the world as a whole, and further assumes an objective relationship between the terms of a language and the sentences in that language, and then claims a reference relationship between the parts of the world and the terms of the language, needs just model theory to make the relationship between the two levels precise with a minimum of additional assumptions.

But what about our earlier claim that Lewis did not need models to explicate his theory? Does he, or doesn't he have models now? The short answer is:

He does have models, a fact which further underscores Lakoff's observation that model-theoretic semantics is just the symbol theory of meaning made precise. To prove this point, we need only turn to his definition of properties.

While it was possible for him to characterize worlds simply as mereological sums of all their parts, he cannot follow the same method with properties. Although he wants to make the point that properties are objectively existent in the worlds, mereology will not help him with the definition. To try to define a property as the merelogical sum of its instances would be to invite severe problems, as Peter van Inwagen demonstrated:

> One cannot say that a property is a sum of objects and that an object exemplifies a property if it is a part of that property. If sphericality were the sum of all spheres (including the C-nonactual ones) and if exemplification were parthood, then the sum of two spheres would be spherical, as would a cubical part of a sphere. The problem of representing properties and exemplification in mereological terms has no obvious solution.[9]

Lewis, we see, is committed to set theory in his exposition of modal realism. And he would of course readily admit this. Properties, we remember, are for Lewis nothing but maximal sets of objects. And so are propositions. In other words, Lewis enjoins us to forget any notions about

9. van Inwagen, 1986, p.210

the nature of properties and propositions we might have had previous to becoming acquainted with his theory, and to toss them out as so much naive philosophy. In the cold hard world of facts, he informs us, there are no properties and propositions as we thought we knew them, only objectively existing sets of things, at least as concrete as things in this world. The immediate problem we have with this view is its lack of intuitive obviousness. It is not as if Lewis' explanation of the nature of properties and propositions caused the scales of naivete to fall from our eyes, so that we now behold the sets of things that were there all the time. On the contrary, Lewis wishes to produce an analysis of modality coupled with a consistent semantics for our modal language . In the absence of any theory which would be intuitively obvious and uncontroversial to the majority of thinkers, he proposes to model everything dear to us in modal talk in terms of sets of entities and constructs out of these sets. Initially, this is neither intuitive nor uncontroversial. The enterprise gains what plausibility it has from our willingness to posit his set-theoretic constructs as the intended model in a model-theoretic semantic, in order to examine the fit between the model and the language. Admittedly, this fit is impressively good. But it is a fit between a structure of sets and a formal language. The fact that he has replaced propositions with sets of worlds does not help us in bridging the semantic

gap. It merely tells us what sort of entities count as admissible according to his model of what there is.

It is Lewis' own explanation which shows perhaps most clearly that we cannot interpret him as trying to do away with models in favour of some 'direct' reference between symbols and the world. Consider his explanation of properties:

> We have frequent need, in one connection of other, to quantify over properties. If we believe in possible worlds and individuals, and if we believe in set-theoretic constructions out of things we believe in, then we have entities suited to play the role of properties.The simplest plan is to take a property as the set of all its instances - all of them, this - and otherworldly alike. Thus the property of being a donkey comes out as the set of all donkeys, the donkeys of other worlds along with the donkeys of ours (emphasis mine).[10]

This sounds definitive and precise, although the word 'role' should perhaps have caused us to pause. But a little later on Lewis makes his position so abundantly clear that one wonders how anyone could have doubted that he intends to model modal talk using set theory. We read:

> To deserve the name of 'property' is to be suited to play the right theoretical role; or better, to be one of a class of entities which together are suited to play the right role collectively. But it is wrong to speak of the role associated with the word 'property' as if it were fully and uncontroversially settled.... The question worth asking is: which entities, if any, among those we should believe in, can occupy which versions of the property role? My answer is, in part, that sets of possibilia are entities we should believe in which are just right for one version of the property role.

---

10. Lewis, 1986, p.50

> There's no point in insisting that this one is the only rightful conception of the properties.[11]

We need only pay attention to the words Lewis uses in this quote to see what is happening. He talks of versions of a role, and of entities suited to play that role in that version. He does not claim that the properties are out there and need only be ostended and assigned suitable symbols to facilitate reference. Alvin Plantinga, I think, is right when he says:

> There they are: all those different models in which different things play the role of a given proposition or property. We are not to ask which really is that proposition or property; no model is endorsed to the exclusion of the others; all are acceptable, although some are more suitable for some purposes than for others.[12]

Lewis, I take it, starts with an interpretation, or role, of 'property' and 'proposition', and then adopts a model of reality which allows him to select entities to correspond to these roles out of the totality of candidates in the worlds he postulates. The entities are chosen according to their fit with modal language terms. This is fundamentally a model-theoretic approach to semantics. I therefore take it to be established that Lewis' modal realism cannot be considered somehow safe from Putnam's conclusions concerning the indeterminacy of reference.

---

11. Lewis, 1986, pp.55-56
12. Plantinga, 1987, p.217

3.2. Magical Relations

In the preceding chapter we saw that the difficulties concerning reference e which Lewis observed in the ersatzist program have their roots in a much broader range of problems, affecting not only the ersatzer's attempts to explain how it is that his symbols refer to things in the world, but Lewis's own as well. The cause for the difficulties of all metaphysical realists was seen to be an adherence to the attempt of explaining meaning by way of a model-theoretic account of truth.

It is obvious from any but the most cursory examination of the literature on this question that metaphysical realists typically refuse to accept the claimed force of 'Putnam's bomb'. Instead they seem to try to make us understand that Putnam and philosophers who think like him fail to see something vital about the world we live in, and about the way we interact with that world. What exactly this something is, remains somewhat difficult to discern. Putnam, at the 1982 conference of the Eastern Division of the American Philosophical Association in Baltimore, referred to this nysterious 'something' as "spooky" and "medieval sounding". But maybe it is, at bottom, nothing more than the feeling that there has to be something to constrain ramp[ant indeterminacy, paired with the realist intuition that this 'something' is an external constraint imposed by objective reality on all our reasoning. When Saul Kripke names causal

chains leading from the external world to the symbol strings as productive of referential success, when David Lewis opts instead for elite properties to ensure the same result, they both want to assure us that, somehow, the world just is such that certain referents are eligible, not that they merely have a theory about their eligibility which is satisfied in some model.

Philosophers opposed to the view which postulates the existence of 'objective' similarity, 'elite' properties, or classifications independent of human classificatory activity most often profess not to understand such talk. What does it mean, they ask, to talk about something existing without presenting the mind with an interpretation in terms of which to understand the existents?

I don't think that much is to be gained from such confrontations. They seem for the most part to be occasioned by incommensurable guiding intuitions on both sides. However, I do believe that Putnam was right in abandoning the attempt to achieve a theory of meaning by way of a model-theoretic account of truth with some added constraints. As I have tried to demonstrate, if his results apply to any metaphysical realist position, they apply to all. To show that this is indeed so, we need not feel condemned to the Sisyphus-labour of producing a conclusive and all-persuasive argument in favour of the anti-realist

position. The metaphysical realist enterprise is imploding under the impact of its own internal divisiveness.

3.2.1. Truth Value Selection

Let us recall Lewis's charge against the magical ersatzer's attempt to take his 'states of affairs' as primitive simples which are selected 'true' or 'false' by the concrete world as the case may be. Lewis showed that the ersatzer not only has to take the relation of truth value selection which the concrete world bears to states of affairs as primitive, a move which Lewis is willing to allow (after all, we do have to take some things as primitive in our theories), he also cannot even tell us what kind of a relation this is supposed to be. It cannot be an external relation, this much seems clear, and it cannot be an internal relation. Lewis professes not to be able to imagine how it might be possible, even in principle, to define something as a relation which is neither external nor internal, and concludes that such a relation can only be grasped by magic.

He contrasts such a rather obscurantist approach to model theoretic semantics with his own, which spells out the truth value selection relation in admirable detail. For Lewis, a proposition is simply a set of possible worlds. The set of all and only those worlds which include a flying pig as at least one of their parts are the proposition, or the state or affairs, or the possibility, or the way things

might be of 'there being a flying pig', or 'that there is a flying pig'. What makes such a proposition true is always relative to a given concrete world. The proposition is true in a world if and only if the proposition has that world as one of its members. Thus, the proposition 'that there is a flying pig' is false in our world, because our world is not a member of the set of worlds which have a flying pig as a part.

For Lewis, propositions, or states of affairs, can be defined without saying anything about their being true or false. And what makes any given proposition true or false can be defined in terms of three well understood relations:

1. the spatiotemporal relation
2. the relation of being part of
3. the relation of being a member of

According to Peter van Inwagen, it was Alvin Plantinga who first realized that not all was as well as it might be with Lewis's new definition of propositions. And although Plantinga himself did not further investigate his suspicions, he suggested his thoughts to van Inwagen who elaborated on them and published the results in 1986.[13]

Simply stated, Plantinga and van Inwagen note that, if they should be forced to admit that they cannot explain what sort of relation their truth value selection relation is,

---

13. van Inwagen relates this in a footnote to his (1986), p.213

then neither can Lewis explain what sort of relation set membership might be. And here is why:

Given three individuals, X,Y, and Z, and three further objects {X,Y}, {X,Z}, and {Y,Z}, what does it mean to say that a given individual bears 'is a member of' to one of the other objects? As Lewis does in his own arguments against the ersatzist position, Plantinga and van Inwagen do not ask for an analysis of that relation. They grant Lewis no less than he is willing to grant them. So Lewis is free to declare that the relation of set membership is one of his primitives. But, they say, he must at least be able to explain whether the relation is an internal one or an external one - just the demand he had placed on their own theories earlier on. And such an explanation seems all the more necessary in light of the fact that we cannot seem to individuate sets by any other means than by specifying the membership relation.

To pick out the set {X,Y} uniquely, we must specify the individuals which bear 'being a member of' to that set, and it must be the case that nothing else bears that relation to that set. Does this mean that membership is an internal relation? Not internal simpliciter, it would seem, at least not if we share Lewis's belief in the possibility of distinct individuals with identical intrinsic properties (as Lewis concedes). For suppose X,Y, and Z were three distinct individuals, and that X and Y had the same intrinsic

properties. Then it seems perfectly correct to say that X is a member of {X,Z}, but that Y is not a member of that set.

So maybe the relation in question is somehow unique in its structure. Maybe individuals with the same intrinsic properties can bear this relation to different sets; but if an individual bears it to any set, it also bears it to all sets with the same intrinsic properties as that set. Let us call such a relation a s-internal relation.

First of all, of course, this obliges us to come up with some plausible candidates for the role of intrinsic property in the case of sets. The property of being a set will not be of much use here, since it does not allow us to individuate between sets. As van Inwagen notes, there seem to be only two types of properties which we might fruitfully consider for this role. They are:

1. the property of containing a given object or objects.
2. the property of containing various numbers of objects having various intrinsic properties.

To flesh this out a bit, let us consider the set {van Inwagen, Lewis}. Given the first type of property as intrinsic, we would define one of the intrinsic properties of our set as 'having Lewis as a member'. Given the second type of property as intrinsic, we would obtain a listing like: 'having at least one human member', 'having no non-human members', or some such. Properties like these may not be universally acceptable as intrinsic properties of sets,

but at least one thing might be said in their favour: they differ from such extrinsic properties of the set (van Inwagen, Lewis) as 'having a member who believes in a plurality of worlds', or 'being equinumerous to the hydrogen atoms in a water molecule' by being less obviously extrinsic.

So let us give Lewis as much as we can and consider all possible cases. We begin with the assumption that either the property of 'containing a given object or objects' uniquely picks out a set, or the two properties of 'containing a given object or objects' and 'containing various numbers of objects having certain intrinsic properties' together uniquely pick out a set. In order to understand this better, it may help to return for a moment to the quasi-Platonic world of section 2, which served us so well in appreciating the force of Lewis's objection to magical ersatzism. Again we encounter a world in which objects bear a very special relationship to entities in a realm which we have no direct access to. As before, we cannot simply behold these 'forms', and distinguish them one from the other. Our only way of knowing them is by way of the special relationship they bear to certain individuals. However, this relationship is not exactly the same as in our previous example. To individuals with the same intrinsic properties can bear it to different 'forms', but if an individual bears it to any 'form', it also bears it to all other 'forms' having the same intrinsic

properties as that 'form'. We might call this relation 's-participation'.

How would we determine if this relation is a s-internal one? By showing that the relation is a consequence of the intrinsic properties of the relata. So, which are the relata? In our example we have on the one hand the two individuals, van Inwagen and Lewis, and on the other hand the set containing them as sole members. We can individuate van Inwagen and Lewis independently of their membership in the set {van Inwagen, Lewis}, and determine what, if any, their intrinsic properties are. But what about the set containing them as its members? The only relevant information about its intrinsic properties is the fact that it bears the s-participation relation to the individuals van Inwagen and Lewis. But the s-participation relation is just what we are trying to understand in terms of the relata. Therefore, if we cannot individuate the things which the s-participation relation is borne to, except by invoking the relation itself in the explanation, then we cannot claim to understand that relation. Again, it would take considerable magical powers to understand a relation between two relata, one of which we cannot individuate independently of that relation.

Applied to our problem of explaining the nature of the set membership relation, if it is internal at all, either

simpliciter, or in the qualified sense we gave it in the foregoing example, then we do not understand it at all.

Maybe, then, we should consider set membership an external relation. But again, Lewis's own reasoning contra the magical ersatzers serves to undermine any such attempt. In the magical ersatzer's case he had argued - correctly, I believe - that if the truth value selection relation was external, then it did not matter what intrinsic properties the relata might have, and consequently it should be possible to explain the relation even if we assumed that the relata had no intrinsic properties whatever.

Applied to his own set-theoretical constructs, if set membership is indeed an external relation, then we should be able to explain it while assuming that the sets have no intrinsic properties at all. To adapt a van Inwagen example to our case, if there is nothing about the intrinsic properties of {van Inwagen, Lewis}, and {van Inwagen, Plantinga} which plays a role in determining set membership, if the relation of 'being a member of' is really an external one, then why couldn't Plantinga bear this relation as easily to {van Inwagen, Lewis} as to {van Inwagen, Plantinga}? On this view, although the intrinsic properties of sets play no role in determining set membership, set membership does seem to depend on the existence of the set's members. Necessarily, it would appear, if van Inwagen and Lewis exist, they belong to {van Inwagen, Lewis}. Van

Inwagen wants to know how this could be so. If set membership is an external relation, then the fact that Plantinga exists, and the fact that he belongs to {van Inwagen, Plantinga}, must be considered unrelated. So what is there to prevent Plantinga from being a member of {van Inwagen, Lewis}? Why can't anything belong to any set? But this is clearly absurd. Plantinga could not belong to {van Inwagen, Lewis}, and that is all there is to it. Which is to say that set membership is not external, and therefore we do not understand it.

This is a result which is not nearly as surprising in the [illegible]tnam's discussion of the shortcomings of a m[illegible] approach to semantic theory as it might have [illegible]. After all, Lewis wants us to see [illegible]lied in set-theoretic terms. In other [illegible]s that the symbol strings of languages [illegible] set theoretic structures, and he subsequently fills out these structures with entities suitable to play the role. As we noted before, he did not have to model propositions in set-theoretical terms. It would have been sufficient to consider propositions to be mereological sums of worlds. What this means is that there is nothing 'out there' which makes propositions be sets rather than mereological sums. And that means that there is nothing in the intrinsic properties of the parts of propositions which makes them members of sets rather than parts of sums.

Consequently, the set membership relation cannot be considered necessary for propositions.

As with all relations connecting 'things in the world' to symbolic representations of them, no matter if we explicitly use model-theoretic interfaces or not, it remains unexplainable how these relationships work, despite any strong intuitions urging us to believe in their existence and efficacy. This often goes unnoticed in the replies of thoroughgoing realists to attacks on their positions, because they can claim to have access or insight into some matters of fact which their opponents either misconstrue or fail to see altogether. Lewis, however, in attempting to construct a theory of modality free of the weaknesses he perceived in previous attempts, has inadvertently built the bomb with which to shake not only the ersatzer's positions, but his own as well.

3.3 The Principle of Recombination

We recall that in section 1.8 of OPW, Lewis had initially attempted to convey the abundance of possible worlds and the fact that there are no gaps in "logical space" by stating that

> 1. absolutely every way that a world could possibly be is a way that some world is, and
>
> 2. absolutely every way that a part of a world could possibly be is the way that some part of some world is.[15]

However, this particular formulation of his theory soon came under attack from many modal theorists. First of all they criticized his decision to identify 'ways things are' with the 'things' themselves. To quote Stalnaker:

> The way things are is a property or a state of the world, not the world itself. The statement that the world is the way it is true in a sense, but not when read as an identity statement (Compare: "the way the world is is the world").[16]

This insistence on retaining the distinction between the property of being a particular thing and the thing itself was echoed elsewhere and seems well worth examining. However, Lewis himself decided to abandon the definitions of abundance and completeness with respect to possible worlds under the impact of another, although related, observation which Peter van Inwagen brought to his attention. Briefly stated, what van Inwagen pointed out to Lewis was that

---

15. Lewis, 1986, p.86
16. Robert Stalnaker, 1976, **Possible Worlds**, Nous 10, p.68

statements (1) and (2) were contentless under any interpretation of its elements which Lewis might be prepared to accept. For to state that every that every way a world could possibly be is a way that some world is just amounts to stating trivially that every world is identical to some world. This does not tell us anything about how many of them there might be. Maybe there is indeed only one world. Or maybe there are exactly 23 worlds. Lewis's principle of plentitude, which he hoped would express the abundance of worlds, would be true for any number of worlds.

This does not mean that statements (1) and (2) should be eliminated from Lewis's modal theory. He quite obviously needs those definitions in order to provide a reductive analysis of modality in nonmodal terms. But he has to look elsewhere for support to his claim that there are no gaps in logical space, that the worlds are in fact abundant enough to exhaust all our desires for possibilia.

3.3.1 The Worlds According to Hume

The place Lewis chooses to look for a way out of this predicament is unexpected, to say the least. He proposes that we look to David Hume's **An Enquiry Concerning Human Understanding** to find the needed possibilia-generating principles.

> To which end, I suggest that we look to the Humean denial of necessary connections between distinct existences. To express the plentitude of possible worlds, I require a principle of recombination according to which patching together parts of different

> possible worlds yields another possible world. Roughly speaking, the principle is that anything can coexist with anything else, at least provided they occupy distinct spatiotemporal positions.[17]

To say, as Lewis does, that he has taken the Humean thesis about laws and causation and applied it to possibility ("Same thesis, different emphasis"[18]), should give even the most sympathetic of his readers some pause. But for now, let us assume the principle of recombination to be a philosophically innocent rendering of certain matters of fact in the vast and wondrous space of possibilia. For even if we are prepared to accept it uncritically and to generate our ideas of modality in accordance with it, we encounter grave problems.

P. Forrest and D.M. Armstrong brought this to our attention in an article entitled: **An Argument Against David Lewis's Theory of Possible Worlds.**[19] They argued that Lewis's principle of recombination allows more possibilia than there can be. It is a necessary consequence of the fact that logical space is complete that there are no vacancies in it. There are no gaps where a possible world should have been but isn't. Likewise, there is no possible world that does not have a place within logical space. What is possible is a matter of fact in some world. Possible worlds exist simpliciter and prior to our inquiry concerning them.

---

17. Lewis, 1986, pp.87-88
18. Lewis, 1986, p.91
19. P. Forrest & D.M. Armstrong, 1984, **An Argument Against David Lewis's Theory of Possible Worlds**, Australasian Journal of Philosophy, 62, pp.164-168

According to Lewis's thesis, for any two worlds in the totality of worlds, there will also be a world consisting of a combination of exactly one duplicate of each of these worlds. This is because each possible world is a possible individual in its own right and may be combined with another distinct individual to form a new one.

But what if we take all the possible worlds there are and combine them to obtain one big world, which by definition contains duplicates of all the possible worlds there are? If the big world is a possible world, it must have already been a member of the original set.

It seems uncontroversial to determine the size of a given world by the cardinality of the set of electrons it contains. Given that the big world does in fact contain a determinate number of electrons, and that it contains a duplicate of itself as a proper part, the addition of the numbers of electrons from the duplicates of all the other worlds, which are also proper parts of the big world would make the big world bigger than it is. We have a reductio.

How does Lewis propose to avoid this unwelcome result? He introduces a proviso qualifying the allowable sizes of spacetime: "size and shape permitting"[20]. What he means by this is not easily understood. In fact, Forrest and Armstrong flatly reject it as ad hoc. Lewis does nothing to dispel this uncomfortable feeling. Instead he shares with us

20. Lewis, 1986, p.101

a fervent hope of his to the effect that somehow or other there just exists a natural(sic) break in the hierarchy of mathematical isomorphs of spacetime manifolds, such that it would allow us to reject certain spacetime sizes as impossible.

Now, we might indeed feel inclined to allow the possibility that the size of spacetime is limited somehow, such that certain claims about it could be considered false. But surely, more than isomorphism between spacetime itself and certain mathematical structures would be needed in order for us to base our acceptance or rejection of particular spacetime sizes on the existence of breaks in mathematical structures. What, after all, could it mean to speak of 'natural' breaks in mathematical structures? Or even of 'breaks' of whatever kind giving us reason to consider the structures occurring below as legitimate isomorphs of existing spacetimes, but the structures above as illegitimate on the grounds that they are not in fact isomorphic to anything? Remember that the question cannot be if they could be isomorphs of existing sizes, but rather if they in fact are. For, according to David Lewis, talk about modality should be understood as talking nonmodally about concrete worlds within logical space. Consequently, we cannot speak about the size of a possible world in modal terms. Possible worlds have concrete sizes, and that is what makes that size 'possible' in our own 'actual' world.

So let us assume that the study of the mathematical structures which we are prepared to consider isomorphs of possible spacetimes actually revealed whatever 'breaks' might be considered to be. And to make this point a little more obvious, let us also assume that the study of the mathematical structures we are prepared to accept as isomorphic with respect to the size and shape of particular crystalline lattices on certain possible worlds reveals certain breaks. All these sizes are matters of fact. In the totality of possible worlds, there is a totality of crystalline lattices which exist in these possible worlds, and this totality is neither more nor less what it really is. For all we know, then, and before a modal theorist's thesis allows me to separate my true modal knowledge from my false modal knowledge, there might be no crystalline lattices at all (prior to lending our ear to David Lewis, it might have seemed at least naively possible to suppose that there are certain kinds of crystalline lattices which do not exist in any possible world whatsoever). Lewis, of course, wants to deny this, and marshals an impressive arsenal of theoretical devices in order to substantiate his view that whenever we speak of possibility, we are actually referring to a fact which is the case at a world spatiotemporally and causally isolated from ours. But it is an integral part of his theory that all these worlds exist independently and antecedently to anything we might find out about them. As he

puts it in OPW on page 104, all the worlds, "no matter how many there may be, get in already on the ground floor."

As we saw at the beginning of this section, however, there is some trouble trying to come up with a definition of the abundance of possible worlds which does not employ modal notions and thus become trivial. We had to abandon the hope that 'absolutely every way that a world could possibly be is a way that some world is' would be useful in defining for us the structure and variegatedness of entities in logical space. This would be true even if there was only one world. Consequently, Lewis introduces the 'principle of recombination: anything can coexist with anything else'. But actually, this definition should have been unacceptable to us the first time we encountered it. We should have asked what the world 'can' is doing in a reductive analysis of modality which aims at expressing modal facts in nonmodal terms. Lewis's principle should have read:

Everything does coexist with everything else.

And if, as we would expect, this would seem too strong a claim to make for Lewis, considering the kind of argument Forrest and Armstrong have advanced against his view, then we would have to qualify either the scope (no longer including everything), or redefine coexistence (so as to exclude certain existents from being able to bear this relation to certain other existents). Since 'coexistence' has previously been defined as 'being spatiotemporally

related', or at least 'being related in a way analogous to spacetime-relatedness' (presumably such that any mathematical structure considered isomorphic to spacetime would also be considered isomorphic to quasi-spacetime), and since this is an external relation, it seems implausible to try to base any restrictions on the intrinsic properties of the relata. In and of itself, then, everything remains admissible for inclusion in the principle of recombination. And, of course, the principle is not to be understood as descriptive of some concrete process taking place in logical space. Logical space is what it is simpliciter (recall that all the possible worlds get in already on the ground floor). So when Lewis says that duplicates of possible worlds coexist with each other 'size and shape permitting', we will have to unpack this statement in the light of his own theory.

First of all, 'being made up of spatiotemporally related parts' is part of the definition of what a possible world is. Secondly, to coexist means to be spatiotemporally related (disregarding for now relatedness analogous to spacetime relatedness). Thirdly, 'size and shape permitting' just means that certain spacetime-relations exist and certain other do not exist. Consequently, Lewis's unpacked principle of recombination should be read as follows:

> An aggregate of spatiotemporally related individuals is spatiotemporally related to another aggregate of spatiotemporally related individuals, if this particular spatiotemporal relationship between the two aggregates is in fact instantiated.

It is absolutely beside the point if the study of mathematical structures will reveal to us any 'breaks'. The principle of recombination in this form goes as trivial as the 'ways things are' definition before it.

When Forrest and Armstrong complain that Lewis's constraint on the possible size of spacetime (which is naive modal talk meaning the factually existent spacetime sizes in logical space) seems to be ad hoc, they don't mean to say that there are varying degrees of ad-hocness to choose from in this case, as Lewis seems to want us to believe. As always we will have to remember that Lewis wants to produce an analysis of modality which will not need to import primitive modality into the definition. According to the theory he advances, there exists a given population of worlds such that the modal operators 'possibly' and 'necessarily' can be replaced by formulations quantifying only over these worlds in nonmodal terms. The term 'possible world' can easily be misleading in discussions of Lewis's theory, because unlike many other modal theorists, the term 'possible' serves in his theory solely to distinguish the set of all non-actual worlds from the one actual world, where 'actual' is understood as an indexical meaning simply 'this world of ours'.

Lewis believes that there is a fact of the matter about the number and size of the worlds there are. Certain sizes, such as the big world Forrest and Armstrong talk about, are

not instantiated. They are not ruled out by some logical principle which is part of Lewis's theory. Lewis endorses set theory as providing the structure of logical space, a it is set theory which produces the 'big world' which is bigger than itself, given Lewis's claim that there is a fixed totality of worlds. Lewis wants us to accept that there are worlds which do not exist simpliciter, just as there are all the worlds which do exist simpliciter. But while we might feel inclined to go along with him on thi claim, we would be well advised to ask which method we should apply in order to arrive at a sorting procedure w would separate the existing worlds from the non-existing ones. As we have seen, his principle of recombination ei leads to a reductio, or, with the necessary proviso adde just states that all and only the spatiotemporally relat sums of individuals are instantiated.

This is what I take to be fundamentally behind the criticism of Forrest and Armstrong. In my opinion, they correctly reformulated Lewis's proviso of 'size and shap permitting' to read 'provided it does not exceed the existing maximal size for worlds', and concluded that th was not a proviso which would eliminate their big world, unless some sort of definition of maximal size was added it. Since nothing in Lewis's theory of modality seems capable of defining such a maximal size, Forrest and Armstrong considered any intimation of greater or lesse

plausibility for any number of size-candidates purely ad hoc, and correctly so.

Consider what Lewis has to say in defense of his proviso. He wants us to agree with him that there are restrictions on the factual occurrence of spacetime sizes which are more or less arbitrary. He states:

> A restriction to four-dimensional, or to seventeen-dimensional, manifolds looks badly arbitrary; a restriction to finite dimensional manifolds looks much more tolerable.[21]

But why should this be the case? To return once again to our example of the sizes and shapes of crystalline lattices, what restriction might we be justified to impose on their number or size prior to any empirical investigation? Of course I want to claim that there is a cardinality which corresponds to the totality of crystalline lattices of particular sizes and shapes in all possible worlds, but that is a consequence of my adherence to the modal theory of David Lewis, which states that there is in fact a cardinality of possible worlds. But this is a claim, and as long as set theory is my tool for the generation of populations of crystalline lattices in possible worlds, this claim will remain unsubstantiated, because set theory does not instruct us to consider a population of 15 crystalline lattices as less plausible than a population of continuum many. Set theory itself, in fact, does not make any empirical existence claims whatsoever - although there might

21. Lewis, 1986, p.103

be independent considerations which, together with set theory, would do so. The problem is that the totality of worlds with the sole exception of the one actual one are causally and spatiotemporally isolated from us, so that whatever independent reasons there might be will forever remain inaccessible to us. It is true, David Lewis's theory bifurcates knowledge into two different kinds, distinguished by the way we attain it. Thus, it is not unconditionally necessary that we be causally acquainted with the subject matter in order to gain knowledge about it. In the modal case, he says, we attain knowledge by being prepared to expand our existential beliefs for the sake of theoretical unity. Presumably, this is the reason behind the size and shape constraint he proposes for possible worlds. And why should we not go along with him? But then we are forced to notice that the case of theoretical unity would be served by a very, very large number of specific spacetime-size constraints, beginning with the proviso that there is in fact only one possible world, namely this one in which we live. Surely, no one would want to doubt that such a theory would provide us with considerable theoretical unity, albeit at the cost of consigning a great number of entities we thought existed to the null set. It is, after all, only the limiting case of Lewis's own brand of modal realism.

Well, maybe what Lewis would actually like to achieve is a theoretical unity, elegance and simplicity which preserves as much of what we pretheoretically consider to be

our modal knowledge as possible. But I do not see how this could be the case. Surely, most of us in our naive pretheoretical stage before the advent of modal realism actually thought that statements about what is possible for me were independent of facts about people in spatiotemporally isolated worlds. And when I claim the knowledge that it is possible that it could have snowed yesterday, I believe that to be a knowledge claim concerning the world we live in, and even were I to believe in a plurality of spatiotemporally isolated universes, in none of which it really did snow yesterday, this would not cause me to revise my knowledge claim concerning my own world in the least. All this spoken naively, of course.

So it is not my or anyone else's pre-analyzed modal knowledge Lewis would like to preserve. In fact, he thinks that anyone who is not a modal realist is simply mistaken about what he actually knows, just attempting some naive interpretations of a phenomenon beyond his grasp, as it were. What Lewis would like to preserve are the modal statements we make and the truth conditions we assign them.

3.3.2 The Problem With The Intended Interpretation

Lewis is fond of stressing the parallels between his account of modal knowledge and some (to him most plausible) accounts of mathematical knowledge. But in light of the Forrest & Armstrong results it seems necessary to point out some obvious differences which should make us think.

Let us imagine that Lewis had been a mathematical theorist who advanced the claim that there was an existent totality of mathematical objects whose organizational principle was the principle of recombination. Let's say that he had come up with a name for the various combinations of mathematical objects his principle generates. He calls them 'Lewis-sets'. Let us further imagine that the relationship between the objects which recombine to form Lewis-sets was called 'Lewis-membership'.

Obviously, this particular 'Lewis-set theory' does not bear much resemblance to the sort of set theory we are acquainted with. But there should be just enough similarity to illustrate the problem with Lewis's view I want to discuss here. Let us suppose Forrest and Armstrong had chosen to take issue with this new and strange theory of mathematical objects and their structure. They perform their previously mentioned thought experiment, not on the possible worlds of modal realism, but on the mathematical entities called Lewis-sets instead. It turns out that there seems to one particular object, or Lewis-set, namely the one having the totality of all Lewis-sets as its proper parts, which appears to be bigger than itself.

Now, suppose Lewis were unhappy with this result, as we might well imagine him to be. So he comes up with a proviso to his principle which is designed to avoid such unwelcome consequences as these. The proviso is: the size of Lewis-

sets permitting. But the size of Lewis-sets had just been defined in terms of the Lewis-membership relation which in turn was a specification of the principle of recombination. Specifically, Lewis-sets were defined in terms of their members, thereby acknowledging the ability of individuals contained in the Lewis-sets to bear this relationship to one another. Therefore, Lewis would have to admit that his proviso fails to avoid the Lewis-set which is bigger than itself.

His dilemma in the case of modal theory is no easier to solve. When he speaks of 'spacetime', he must make it clear to us what he takes this concept to mean. On the interpretation which we at first take Lewis to provide in section 1.6 of OPW, spacetime is a name for a specific relation between objects. According to this view, worlds are mereological sums of spatiotemporally related things, a view which, as Lewis clearly states, "makes no provision for an absolutely empty world."[22]

This much seems clear, then: worlds just are the sums of the things they contain, and what separates one world from another is that only mereological sums of things bearing the spacetime relation to each other qualify as distinct worlds.

But six lines down from the quote just given, we read: "There can be nothing much: just some homogeneous unoccupied

22. Lewis, 1986, p.73

spacetime, or maybe only one single point of it."[23] And this should confuse us. For spacetime is either a relation, or it is a thing. According to Lewis, things and the relations between them are all the elements necessary to account for the totality of possibilia (with the addition of properties, of course). But now it seems as if we had an odd crossbreed on our hands which has the characteristics of both. Does this make sense? I do not think so.

To make this more apparent, consider a world, one of Lewis's minimal ones, which consisted only of some unoccupied spacetime. In other words, we now have it that there can be space without anything occupying it. Space simpliciter. And space exists, in fact it is something which can be counted among the things which form the mereological sum called a 'world'. Likewise with time. It also exists simpliciter. And furthermore, the two existents enter into a relationship with each other, called the spatiotemporal relationship, which generates spacetime points. I take it that it is this idea which makes Lewis balk at Forrest and Armstrong's conclusions. This is why he thinks that there can be a study of 'spacetime' which will reveal the extant sizes of it, and therefore the maximal number of individuals which come together to form worlds. But this is absurd.

If spacetime points are existents which can, by themselves, form worlds of 'unoccupied spacetime', as Lewis

---

23. Lewis, 1986, p.73

clearly presupposes, then they must be counted as parts of the mereological objects which worlds are. If we go further along with Lewis to suppose that it is the spacetime relation which unifies worlds (again omitting, for the sake of simplicity, any stranger relations of a more alien nature), then spacetime points must bear this relation to one another. Obviously, by the principle of recombination, even with proviso, I can take a minimal world of one single spacetime point and combine it with a slightly less minimal world of three spacetime points to obtain a world containing exactly four spacetime points. The glue for this new world, as for any other, is the spacetime relatedness of its members. Unfortunately, it is wholly unclear how we would spell out the nature of this relatedness, external as it is supposed to be, in the case of spacetime points.

And a further problem: given that we were convinced somehow to accept these minimal worlds as genuine, how would we go about combining them with the other things which make up worlds, such as donkeys and pigs, or protons, or angels? Spacetime points are things which constitute worlds. They are related spatiotemporally, whatever that may turn out to be, and this relation is an external one. Now we want to claim that my favourite cat is related to a specific number of such interrelated points by virtue of that same relation, and furthermore, that a minimal world comprised of a suitably large number of unoccupied spacetime points can be combined with a duplicate of my favourite cat to yield

another world consisting of all and only those unoccupied spacetime points and the duplicate of my cat. What sort of relation might this be, which a spacepoint(?) can bear to a timepoint(?) to become a spacetime point, which spacetime points may bear to each other, which cats may bear to spacetime points, and finally, which cats may bear to other cats, or protons, or angels?

As Lewis explains to us, "what we call'spatiotemporal relations' are relations that behave in the relativistic way, with spatial or temporal distance, but not both." By contrast, for a Newtonian world, the "theory will say that any two spacetime points are related by a spatial distance and a temporal distance."[24] So we have a pretty good idea what to understand by spatiotemporal relatedness, even if the worlds should be Newtonian rather than relativistic. In either case, the relation is defined in terms of distance, either spatial or temporal or both. We might be prepared to agree that such a relation does in fact exist between things in our world, things like cats and mats and cherries and trees. And let us pretend that Lewis had convinced us that it also held between those things we call spacetime points. But how does it hold between the spacetime points and the more usual sorts of things we might want to combine with them? How do we express the spacetime distance between a number of spacetime points and the thing or things which supposedly occupy them?

---

24. Lewis, 1986, p.74

It should seem obvious that this is an absurd question. If spacetime points do exist such that they can constitute worlds by themselves, then we will need another relation which can define for us how they are connected to things like electrons and donkeys. If, on the other hand, spacetime points are to be considered to be the smallest constituents of what there is in a world, then it does not make sense to talk about unoccupied spacetime points.

In summary, as long as Lewis defines the worldmaking relation as the spatiotemporal one, and does not give us a special, secondary relation to determine how the bare worlds (construed as mereological sums of unoccupied spacetime points) are to be occupied, we do not have a way to speak about the kinds of possible worlds we should be most interested in, namely the occupied ones. If, on the other hand, he intends the spatiotemporal relation to be borne by things to other things, not construed as spacetime points here, then it makes no sense to speak of spacetime size or shape independently of the things bearing the spatiotemporal relation to one another. Thus, in providing the principle of recombination with the proviso 'size and shape permitting', Lewis has to make a choice about how to interpret his worlds. If he wants spacetime points to be able to exist unoccupied, then he has not told us how his theory gives us the occupied worlds, although it may then make sense to wonder about the maximal size and shape of worlds. But if the spatiotemporal relation is meant to unite individuals

other than spacetime points with each other, than we have no grounds to suppose that unoccupied spacetimes are possible, because on that interpretation talk about spacetime points is simply talk about the relational properties of the things occupying them. And consequently, we have no grounds to add provisos such as 'size and shape permitting' to our principle of recombination, because size and shape cannot be defined independently of the individuals constituting the size and the shape.

Lewis had conceded that his theory would be in serious trouble if the study of mathematical generalizations of ordinary spacetime manifolds revealed no suitable breaks indicating worlds which exist and worlds which do not exist. But it seems that he is in trouble regardless of the outcome of any such study, because his proviso is ineffective and mathematical study will therefore not offer any help.

Apart from any criticism which attacks a reductive analysis of modality on the grounds that any interpretation of modal language in terms of existing things fail to have any bearing on the 'intentions' which make human beings use modal language in the first place, and therefore only provides a model of reality which fits modal language as an uninterpreted but nevertheless fully instantiated phenomenon, his theory itself fails to generate the totality of possibilia which he had supposed logical space to be filled with. The sort of primitive modality Lewis had called

"the most obnoxious kind" in his criticism of the ersatzist enterprise infects his own modal realism just as obnoxiously, because, as it turns out, the totality of possible worlds is what it is simply because it 'could' not be anything else.

BIBLIOGRAPHY

Benacerraf, P. & Putnam, H. (ed), 1964, **Philosophy of Mathematics**, Englewood Cliffs, Prentice-Hall Inc.

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

Castongay, C., 1972, **Meaning and Existence in Mathematics**, New York, Springer Verlag

Elgin, C.Z., 1983, **With Reference to Reference**, Indianapolis, Hackett Publishing

Forbes, G., 1982, **Canonical Counterpart Theory**, Analysis 42
----------, 1985, **The Methaphysics of Modality**, Oxford, Clarendon Press

Forrest, P. & Armstrong, D.M., 1984, **An Argument Against David Lewis' Theory of Possible Worlds**, Australasian Journal of Philosophy, Vol.62

Goodman, N., 1978, **Ways of Worldmaking**, Indianapolis, Hackett Publishing Co.
----------, 1984, **Of Mind and Other Matters**, Cambridge, Harvard University Press

Haack, S., 1978, **Philosophy of Logics**, Cambridge, Cambridge University Press

Hughes, G.E. & Cresswell, M.J., 1968, **An Introduction to Modal Logic**, London, Methuen & Co. Ltd.

Kripke, S., 1980, **Naming and Necessity**, Cambridge, Harvard University Press

Lakoff, G., 1987, **Women, Fire, and Dangerous Things**, Chicago, The University of Chicago Press

Lewis, D., 1973, **Counterfactuals**, Oxford, Basil Blackwell Ltd.
---------, 1984, **Putnam's Paradox**, Australasian Journal of Philosophy, Vol.62, No.3
---------, 1986, **On the Plurality of Worlds**, Oxford, Basil Blackwell Ltd.

Plantinga, A., 1974, **The Nature of Necessity**, Oxford, Oxford University Press
----------, 1987, **Two Concepts of Modality**, in Tomberlin, J.E. (ed), Philosophical Perspectives, 1

Putnam, H., 1977, **Realism and Reason**, Proceedings of the American Philosophical Association, 50

----------, 1980, **Models and Reality**, The Journal of Symbolic Logic, Vol.45, No.3

----------, 1981, Reason, **Truth and History**, Cambridge, Cambridge University Press

Quine, W.v.O., 1953, **From a Logical Point of View**, Cambridge, Harvard University Press

Salmon, N., 1981, **Reference and Essence**, Princeton, Princeton University Press

Stalnaker, R.C., 1976, **Possible Worlds**, Nous 10

Tomberlin, J.E. (ed), 1987, **Philosophical Perspectives, 1**, Atascadero, Ridgeview Publishing Co.

van Inwagen, P., 1986, **Two Concepts of Possible Worlds**, Midwest Studies in Philosophy, Vol.XI, Minneapolis, University of Minnesota Press