# Ways a World Might Be
Robert C. Stalnaker

Possible Worlds (1976/1984)

## Abstract and Keywords

This paper explores David Lewis's four theses on possible worlds. It is argued that these constitute a doctrine called extreme realism about possible worlds, which is deemed false. However, these theses need not be accepted or rejected as a package. The independence of the more plausible parts of the package is shown to defend the coherence of a more moderate form of realism about possible worlds, one that may be justified by common modal opinions and defended as a foundation for a theory about the activities of rational agents.

Keywords:  possible worlds, David Lewis, extreme realism, coherence, rational agents

## Abstract and Keywords

This paper explores David Lewis's four theses on possible worlds. It is argued that these constitute a doctrine called extreme realism about possible worlds, which is deemed false. However, these theses need not be accepted or rejected as a package. The independence of the more plausible parts of the package is shown to defend the coherence of a more moderate form of realism about possible worlds, one that may be justified by common modal opinions and defended as a foundation for a theory about the activities of rational agents.

*Keywords:*  possible worlds, David Lewis, extreme realism, coherence, rational agents

According to Leibniz, the universe—the actual world—is one of an infinite number of possible worlds existing in the mind of God. God created the universe by actualizing one of those possible worlds—the best one. It is a striking image, this picture of an infinite swarm of total universes, each by its natural inclination for existence striving for a position that can be occupied by only one, with God, in his infinite wisdom and benevolence, settling the competition by selecting the most worthy candidate.

But in these enlightened times, we find it difficult to take this metaphysical myth any more seriously than the other less abstract creation stories told by our primitive ancestors. Even the more recent expurgated versions of the story, leaving out God and the notoriously chauvinistic thesis that our world is better than all the rest, are generally regarded, at best, as fanciful metaphors for a more sober reality. J. L. Mackie, for example, writes, 'talk of possible worlds...cries out for further analysis. There *are* no possible worlds except the actual one; so what are we up to when we talk about them?' [1] Larry Powers puts the point more bluntly: 'The whole idea of possible worlds (perhaps laid out in space like raisins in a pudding) seems ludicrous.' [2]

These expressions of skepticism and calls for further analysis are of course not directed at Leibniz but at recent uses of parts of his metaphysical myth to motivate and give content to formal semantics for modal logics. In both formal and philosophical discussions of modality, the concept of a possible world has shown itself to have considerable heuristic power. But, critics have argued, a heuristic device should not be confused with an explanation. If analyses of modal concepts (or of the concept of a proposition) in terms of possible worlds are to be more than heuristic aids in mapping the relationships among the formulae of a modal logic, the concept of a possible world itself must be explained and justified.

Although it is commonly taken to be an obvious truth that there really are no such things as possible worlds—that the myth, whether illuminating or misleading, explanatory or obfuscating, is nevertheless a myth—this common opinion can be challenged. That is, one might respond to the possible worlds skeptic not by explaining the metaphor but by taking the (p. 26 ) story to be the literal truth. David Lewis responds in this way, and he cites common opinion and ordinary language on his side:

> I believe there are possible worlds other than the one we happen to inhabit. If an argument is wanted, it is this: It is uncontroversially true that things might have been otherwise than they are. I believe, and so do you, that things could have been different in countless ways. But what does this mean? Ordinary language permits the paraphrase: there are many ways things could have been besides the way that they actually are. On the face of it, this sentence is an existential quantification. It says that there exist many entities of a certain description, to wit, 'ways things could have been'. I believe permissible paraphrases of what I believe; taking the paraphrase at its face value, I therefore believe in the existence of entities which might be called 'ways things could have been.' I prefer to call them 'possible worlds'. [3]

Lewis does not intend this as a knockdown argument. It is only a presumption that the sentences of ordinary language be taken at face value, and the presumption can be

defeated if the naive reading of the sentences leads to problems which can be avoided by an alternative analysis. The aim of the argument is to shift the burden to the skeptic who, if he is to defeat the argument, must point to the problems which commitment to possible worlds creates, and the alternative analysis which avoids those problems. Lewis does not think the skeptic can do either.

The rhetorical force of Lewis's argument is in the suggestion that possible worlds are really not such alien entities as the metaphysical flavor of this name seems to imply. The argument suggests not that ordinary language and our common beliefs commit us to a weighty metaphysical theory, but rather that what appears to be a weighty metaphysical theory is really just some ordinary beliefs by another name. Believing in possible worlds is like speaking prose. We have been doing it all our lives.

But for this to be convincing, the shift from 'ways things might have been' to 'possible worlds' must be an innocent terminological substitution, and I do not believe that, as Lewis develops the concept of a possible world, it is. To argue this point I will state four theses about possible worlds, all defended by Lewis. Together they constitute a doctrine which I will call extreme realism about possible worlds. It is this doctrine against which the skeptic is reacting, and against which, I shall argue, he is justified in reacting. I believe the doctrine is false, but I also believe that one need not accept or reject the theses as a package. The main burden of my argument will be to show the independence of the more plausible parts of the package, and so to defend the coherence of a more moderate form of realism about possible worlds, one that might be justified by our common modal opinions and defended as a foundation for a theory about the activities of rational agents.

(p. 27 ) Here are Lewis's four theses:

> 1. *Possible worlds exist.* Other possible worlds are just as real as the actual world. They may not actually exist, since to actually exist is to exist in the actual world, but they do, nevertheless, exist.
> 2. *Other possible worlds are things of the same sort as the actual world*—'I and all my surroundings'. [4] They differ 'not in kind, but only in what goes on at them. Our actual world is only one world among others. We call it alone actual not because it differs in kind from all the rest, but because it is the world we inhabit.' [5]
> 3. *The indexical analysis of the adjective 'actual' is the correct analysis.* "The inhabitants of other worlds may truly call their own world actual if they mean by 'actual' what we do; for the meaning we give to 'actual' is such that it refers at any world *i* to that world *i* itself. 'Actual' is indexical, like 'I' or 'here' or 'now': it depends for its reference on the circumstances of utterance, to wit, the world where the utterance is located. [6]

4. *Possible worlds cannot be reduced to something more basic.* 'Possible worlds are what they are and not another thing.' It would be a mistake to identify them with some allegedly more respectable entity, for example a set of sentences of some language. Possible worlds are 'respectable entities in their own right'. [7]

The first thesis, by itself, is compatible with Lewis's soothing claim that believing in possible worlds is doing no more than believing that things might have been different in various ways. What is claimed to exist are things which ordinary language calls 'ways things might have been', things that truth is defined relative to, things that our modal idioms may be understood as quantifiers over. But the first thesis says nothing about the nature of the entities that play these roles. It is the second thesis which gives realism about possible worlds its metaphysical bite, since it implies that possible worlds are not shadowy ways things could be, but concrete particulars, or at least entities which are made up of concrete particulars and events. The actual world is 'I and my surroundings'. Other possible worlds are more things like that. Even a philosopher who had no qualms about abstract objects like numbers, properties, states, and kinds might balk at this proliferation of fullblooded universes which seem less real to us than our own only because we have never been there.

The argument Lewis gives for thesis (1), identifying possible worlds with ways things might have been, seems even to be incompatible with his explanation of possible worlds as more things of the same kind as I and all my surroundings. If possible worlds are ways things might have been, then the actual world ought to be *the way things are* rather than *I and all my surroundings. The way things are* is a property or a state of the world, not (p. 28) the world itself. The statement that the world is the way it is is true in a sense, but not when read as an identity statement (compare: 'the way the world is is the world'). This is important, since if properties can exist uninstantiated, then *the way the world is* could exist even if a world that is that way did not. One could accept thesis (1)—that there really are many ways that things could have been—while denying that there exists anything else that is like the actual world.

Does the force of thesis (2) rest, then, on a simple equivocation between 'the actual world' in the sense that is roughly captured in the paraphrase 'I and all my surroundings' and the sense in which it is equivalent to 'the way things are'? In part, I think, but it also has a deeper motivation. One might argue from thesis (3)—the indexical analysis of actuality—to the conclusion that the essential difference between our world and the others is that we are here and not there.

Thesis (3) seems to imply that the actuality of the actual world—the attribute in virtue of which it is actual—is a world#relative attribute. It is an attribute which our world

has relative to itself but which all other worlds have relative to themselves too. But the moderate realist conception of possible worlds seems to conflict with this world# relativity, or contingency, of the property of actuality; if possible worlds are abstract 'ways things could have been', then there is something distinctive about one of them, which constitutes its actuality: it is the one which corresponds to the world itself. Lewis argues that any 'moderate version of modal realism comes to grief in the end' for this reason. Referring to abstract 'ways things could have been' as 'ersatz worlds', he argues that, on the moderate realist account, 'the actualized ersatz world is special, since it alone represents the one concrete world. *And it is special not just from its own standpoint, but from the standpoint of any world*. So it is noncontingently special, since contingency is variation from world to world. But it is part of the theory that the actualized ersatz world is the special one. So it seems to turn out to be a noncontingent matter which of the ersatz worlds is actualized.' [8]

The mistake in this argument, I think, is in the italicized statement. It *is* a special fact about 'the actual world' (the way things are) that it alone corresponds to the one concrete world. But this is a contingent fact, which means that from the standpoint of other worlds, it is not a fact. From the standpoint of a counterfactual possible world, that world has the special property of being the one and only 'way things could have been' that corresponds to the concrete world. But does this not mean that, looking at things from an objective, absolute standpoint, merely possible people and their surroundings are just as real as we and ours? Only if one identifies the objective or absolute standpoint with a neutral standpoint outside of all possible worlds. But there is no such standpoint. The objective, absolute point of view is the (p. 29) view from within the actual world, and it is part of the concept of actuality that this should be so. We can grant that fictional characters are as right, from their points of view, to affirm their full# blooded reality as we are to affirm ours. But their point of view is fictional, and so what is right from it makes no difference as far as reality is concerned.

My point is that the *semantic* thesis that the indexical analysis of 'actual' is correct can be separated from the metaphysical thesis that the actuality of the actual world is nothing more than a relation between it and things existing in it. Just as one could accept the indexical analysis of personal pronouns and be a solipsist, or accept the indexical analysis of tenses and believe that the past exists only as memory and the future only as anticipation, one can accept the indexical analysis of actuality while excluding from one's ontology any universes that *are* the way things might have been.

Let me see if I can use the analogy between the present time and the actual world (an analogy that Lewis obviously has in mind) to make my point a little clearer. Pretend for the moment that we accept an Augustinian thesis about time, the one expressed in the following quotation:

At any rate it is now quite clear that neither the future nor the past really exist. Nor is it right to say there are three times, past, present, and future. Perhaps it would be more correct to say: there are three times, a present of things past, a present of things present, a present of things future. For these three exist in the mind, and I find them nowhere else: the present of things past is memory, the present of things present is sight, the present of things future is expectation. If we are allowed to speak thus, I see and admit that there are three times, that three times truly are.

By all means continue to say that there are three times, past, present, and future; for, though it is incorrect, custom allows it. By all means say it. I do not mind, I neither argue nor object: provided that you understand what you are saying and do not think future or past now exist. There are few things that we phrase properly; most things we phrase badly; but what we are trying to say is understood. [9]

Now suppose we are trying to give a semantic analysis of tenses. Is our pretended metaphysical belief about time compatible with the standard semantic account—the one that says, for example, that at any time $t$ a future tense statement is true at $t$ just in case the corresponding present tense statement is true at some $t$ later than $t$? I think it is, *provided we understand the times in terms of which our past and future tense statements are interpreted as things which exist in the present*. That is exactly how Augustine does understand them.

Just as we needed to distinguish two senses of 'possible world', so the Augustinian theory needs to distinguish two senses of time. There is the sense in which 'neither the future nor the past really exist', and the sense in which, 'if we are allowed to speak thus', we can correctly say 'that three (p. 30) times truly are'. The future and the past (in the latter sense) exist in the present (former sense), which is to say that they exist. So they are available as resources for the interpretation of tensed statements.

Only the present really exists (we are supposing). So the word 'now' picks out the time that is uniquely real (that is uniquely realized). But, one might object, do not speakers in the past, and in the future, who say 'only the present really exists' refer to their own times, and do they not speak the truth? Yes, they *did*, or *will*, refer to their own times, and they *did*, or *will*, speak the truth. But the fact remains (according to our hypothesized metaphysics) that their times do not exist *now*, which is to say that they do not exist at all (except as memory or anticipation). The fact that those times were or will be real does not show that they *are* real.

I do not want to defend the Augustinian theory of time, or even a less idealistic version that substitutes traces and potentialities for memory and anticipation. My only point is that Augustine and his metaphysical opponents who believe that other times are as real as our own might share the same semantics for tenses. The correctness of the standard semantics for tenses and temporal indexicals is not, in itself, a refutation of a theory that denies the reality of other times. The thesis that the actual world alone is real is superficially analogous to the Augustinian thesis that the present alone is real—that all else, or all that appears to be else, is an aspect of the present. But the Augustinian thesis has content, and can be coherently denied, because it says something substantive about what alone is real. In effect, it says that the actual world is a moment of time. But the thesis that the actual world alone is real has content only if 'the actual world' means something other than the totality of everything there is, and I do not believe that it does. The thesis that there is no room in reality for other things than the actual world is not, like the Augustinian thesis, based on a restrictive theory of what there is room for in reality, but rather in the metaphysically neutral belief that 'the actual world' is just another name for reality.

The extreme realist will, of course, deny this. He must hold that the indexical 'actual' picks out, not the whole of reality, but only a part of it. But what part? How do I draw the boundaries around that part of reality which is appropriately related to me to be part of my actual world? I do not think the extreme realist can give a satisfactory answer to this question. Presumably, any part of reality that is spatially or causally connected with something in the actual world is itself part of the actual world. No one thinks that other possible worlds are literally 'out there' in space (like raisins in a pudding), or that we might communicate with merely possible beings. But if other possible worlds are causally disconnected from us, how do we know anything about them? If the truth or falsity of our modal claims depends on the existence of things and events which are causally disconnected from us, then even the simplest claims about what is possible are unverifiable speculations.

(p. 31 ) My point is not a verificationist one. I am not bothered by the fact, freely admitted by Lewis, that there may be some modal facts—facts about what is possible— which we can never know. My worry is that I do not see how, on Lewis's account, there can be any other kind of modal fact. I do not see how we could ever have any reason to believe that a proposition we know to be false was nevertheless one that might have been true. As Lewis says in the argument quoted above, 'I believe, and so do you, that things could have been different in countless ways.' But we also believe that this belief is a reasonable one, and not a speculation about what is going on in some place so far away that it is not even part of our universe. It is the reasonableness of modal beliefs which I think Lewis's extreme realism cannot account for.

In the course of defending his version of modal realism, Lewis makes some general remarks about his conception of the methodology of metaphysics. The particular problem with extreme realism that I have just pointed to derives, I think, from a limitation in that methodology. Let me quote some of Lewis's remarks:

> One comes to philosophy already endowed with a stock of opinions. It is not the business of philosophy either to undermine or to justify those preexisting opinions, to any great extent, but only to try to discover ways of expanding them into an orderly system. A metaphysician's analysis... succeeds to the extent that (1) it is systematic, and (2) it respects those of our pre#philosophical opinions to which we are firmly attached. Insofar as it does both better than any alternative we have thought of, we give it credence....
>
> So it is throughout metaphysics; and so it is with my doctrine of realism about possible worlds....Realism about possible worlds is an attempt, the only successful attempt I know of, to systematize preexisting modal opinions. [10]

There is much in this conception of metaphysics that seems to me right, but I think we must ask more of our philosophical theories. They must not only systematize preanalytic opinions, but also help to explain the source of those opinions and their role in our practical activities. We want to know not only what it is to believe that things might have been different in various particular ways, but also why those beliefs make a difference to inhabitants of the actual world. We want to know why it is reasonable to have such beliefs, and why it is reasonable to use them in the way we do. Even if extreme realism were to succeed in providing an elegant systematization of modal opinions, it would, I think, preclude a plausible explanation of the practical significance which we believe those opinions to have.

The moderate realism whose coherence I am trying to defend accepts theses (1) and (3), and rejects thesis (2). What about thesis (4)? If we identify possible worlds with ways things might have been, can we still hold that they are 'respectable entities in their own right', irreducible to anything more fundamental? Some may believe that even if a moderate realism about (p. 32) possible worlds is coherent, one buys the added plausibility which it has over extreme realism only at the cost of incurring the obligation to give some further analysis or explanation of the concept of a possible world. If there were possible worlds, in the sense of other things like me and my total environment, then they could perhaps be referred to and quantified over even if not much more could be said about them than that they are other things like our actual universe. As concrete objects, they can stand alone, if they stand at all. One would need no justification for including them in one's theory of reality other than that they are there. But the moderate

realist believes that the only possible worlds there are—ways things might have been—are (like everything that exists at all) elements of our actual world. They obviously are not concrete objects or situations, but abstract objects whose existence is inferred or abstracted from the activities of rational agents. It is thus not implausible to suppose that their existence is in some sense dependent on, and that their natures must be explained in terms of, those activities. Thus Mackie, discussing possible worlds analyses of conditionals, writes,

> possible situations, or possible worlds, just because they are not
> actual (or may not be actual) do not stand on their own, do not exist
> independently....People can consider possibilities; but the possibilities
> exist only as the contents of such considerings...'possible situations'...
> cry out for some further analysis. This analysis must be given in terms of
> what people do. [11]

The claim that possible worlds are dependent on actual human activities and the demand that possible worlds be explained in terms of such activities are vague, but their general intent and motivation are clear enough. I shall try to respond to this concern by locating claims about possible worlds in a theory of rational activities. But I want to resist a more specific demand with the same motivation: this is the demand, which Robert Adams argues must be met if we are to escape commitment to extreme realism, that 'if there are any true statements in which there are said to be nonactual possible worlds, they must be reducible to statements in which the only things there are said to be are things which are in the actual world, and which are not identical with nonactual possibles.' [12] Unless the reminder that by 'possible world' we mean nothing more than 'way things might have been' counts as such a reduction, I do not see why this should be necessary.

Two problems need to be separated: the first is the general worry that the notion of a possible world is a very obscure notion. How can explanations in terms of possible worlds help us to understand anything unless we are told what possible worlds are —and told in terms which are independent of the notions which possible worlds are intended to explain? The second problem is the specific problem that believing in possible worlds and in the (p. 33 ) indexical analysis of actuality seems to commit one to extreme realism, which is obviously false. Now to point to the difference between a way our world might have been and a world which *is* the way our world might have been, and to make clear that the possible worlds whose existence the theory is committed to are the former kind of thing and not the latter, is to do nothing to solve the first problem; in fact, it may make it more acute since it uses a modal auxiliary to explain what a possible world is. But this simple distinction does, I think, dissolve the second problem, which was the motivation for Adams's demand for an analysis.

Not only is an eliminative reduction of possible worlds not necessary to solve the second problem, it also may not be sufficient to solve the first. I shall argue that the particular reduction that Adams proposes—a reduction of possible worlds to propositions—by itself says nothing that answers the critic who finds the concept of a possible world obscure. His reduction says no more, and in fact says less, about propositions and possible worlds than the reverse analysis that I am defending—the analysis of propositions in terms of possible worlds.

Adams's analysis is this:

> Let us say that a *world#story* is a maximal consistent set of propositions. That is, it is a set which has as its members one member of every pair of mutually contradictory propositions, and which is such that it is possible for all of its members to be true together. The notion of a possible world can be given a contextual analysis in terms of world# stories. [13]

For a proposition to be true in some or all possible worlds is for it to be a member of some or all world#stories. Other statements that seem to be about possible worlds are to be replaced in a similar way by statements about world#stories.

There are three undefined notions used in Adams's reduction of possible worlds: *proposition*, *consistent*, and *contradictory*. What are propositions? Adams leaves this question open for further discussion; he suggests that it might be answered in various different ways. Little is said about them except that they are to be thought of as language#independent abstract objects, presumably the potential objects of speech acts and propositional attitudes. And it is, of course, assumed that they have truth#values.

What is consistency? The notion used in the definition of world#story is a property of *sets* of propositions. Intuitively, a set of propositions is consistent if it is possible for all its members to be true together. We cannot, of course, explain this intuitive idea in terms of the existence of a possible world in which all of them are true. Presumably, consistency will be a primitive notion, but we must be sure that the world#story theory includes some constraints on it which ensure that it has the right properties. Two (p. 34 ) assumptions that are necessary, and that should be uncontroversial, are the following:

> (W1) The set of all true propositions is consistent.
> (W2) Any subset of a consistent set is consistent.
>
> What is a contradictory? This relation between proposition might be defined in terms of consistency as follows: *A* and *B* are contradictories if and only if

{*A*, *B*} is not consistent, and for every consistent set of propositions $\Gamma$ either
$\Gamma$ # {*A*} or $\Gamma$ # {*B*} is consistent. The theory tacitly assumes:
(W3) Every proposition has a contradictory.

These definitions and postulates yield a minimal theory of propositions. It is minimal in that it imposes no structure on the propositions except what is required to define the standard propositional relations such as compatibility, implication, and equivalence, and to ensure that these relations have the right properties. Implication, for example, can be defined in a familiar way: *A implies B* if and only if a set consisting of *A* and a contradictory of *B* is not consistent. (W1) and (W2) ensure that implication, defined this way, has the right properties. For example, it is transitive, reflexive, and preserves truth.

This minimal theory is not, however, sufficient to justify what Adams calls 'the intuitively very plausible thesis that possibility is holistic rather than atomistic, in the sense that what is possible is possible only as part of a possible completely determinate world'. [14] The assumptions we have made so far do not imply that every consistent set of propositions is a subset of a world#story, [15] so to ensure that the theory satisfies the principal thesis that motivates it, we must add the following as an additional postulate:

(W4) Every consistent set is a subset of a maximal consistent set.

It will be useful to compare this reduction of possible worlds to propositions with the competing reduction of propositions to possible worlds. What is at stake in choosing which of these two notions to define in terms of the other? Adams refers to the 'not unfamiliar tradeoff between non#actual possibles and intensions (such as propositions); given either, we may be able to construct the other, or do the work that was supposed to be done by talking about the other'. [16] But the two proposals are not equivalent. Part of what distinguishes them is an elusive question of conceptual priority, but (p. 35) there are also more substantive differences, both in the structure imposed on propositions and possible worlds and in the questions left to be answered by further developments of the respective theories.

One might think that a reduction of possible worlds to propositions was more appropriate than the reverse since propositions are more familiar entities than possible worlds, and analysis should explain the less familiar in terms of the more familiar. However problematic propositions may be, we can still point to particular examples of them (like the proposition that Gerald Ford became President of the United States in 1974), and to facts about them with which we are familiar (for example, that the above# mentioned proposition is believed by Richard Nixon, and that it entails that Gerald Ford became either the President or the Vice#President in 1974). Possible worlds, on the

other hand, are a philosopher's invention and are further removed from the phenomena of speech and thought. So they are more in need of analysis.

I think this reasoning is based on a mistaken assumption about the point and content of a philosophical analysis. An analysis makes a claim about a relation among concepts which, if accepted, can be informative in either direction, or in both directions. It may be as helpful in explaining an obscure concept to reduce other things to it as to reduce it to other things. This would not be true if an analysis were like a stipulative definition of a previously meaningless sign, but normally all the terms involved in a philosophical analysis are terms that we understand, to some extent, prior to the analysis.

What is more important than familiarity for choosing the order of an analysis is the question of structure. Normally, one would expect to analyze the more highly structured set of entities in terms of a less highly structured one so that the analysis might help to explain the structure. Propositions are highly structured. One proposition can be stronger or weaker than another, or equivalent to it, compatible or incompatible with it, related to it as contrary or contradictory. These relations give rise to questions about the identity conditions for propositions, about closure conditions for the set of all propositions, about whether some propositions are constructed out of others. Possible worlds, on the other hand, are relatively unstructured. One may, for some purposes, want to define relations between possible worlds, like resemblance in some respect or other, or accessibility of some kind, but such relations are not essential to the concept of a possible world in the way that relations like entailment and compatibility are essential to the concept of a proposition. A minimal theory of propositions and possible worlds needs an account of the propositional relations, but it can get along without an account of relations between possible worlds.

Both of the theories that I am comparing impose structure on the set of propositions, but in the minimal world#story theory that structure derives from a primitive and unexplained concept of consistency. In contrast, the (p. 36) possible worlds analysis of propositions yields definitions of consistency and the other propositional relations in terms of elementary set#theoretic relations between the sets of possible worlds determined by the propositions. Whether one accepts these definitions as explanations of the propositional relations may depend on whether, in the end, one accepts the concept of a possible world as having intuitive content and an independent role to play in a theory; but even if possible worlds are totally opaque postulated entities, the possible worlds analysis of propositions still imposes a structure on the set of propositions which can be tested against intuitions about propositions, and against theoretical assumptions about the role of propositions in a theory of mental states and linguistic actions.

If we set aside questions of conceptual priority—of which concepts and principles should be primitive and which defined or derived—what are the differences between the two theories we are comparing? The world#story theory is weaker, leaving open questions which are settled by the possible worlds analysis of propositions. The following two theses are consequences of the possible worlds analysis but not of the world#story theory; the first is a closure condition; the second concerns identity conditions.

> (W5) For every set of propositions    there is a proposition A such that    implies A, and A implies every member of   .
> (W6) Equivalent propositions are identical.

Are these consequences of the possible worlds analysis welcome or not? Thesis (W5) seems reasonable on almost any theory of propositions and propositional attitudes. What it says is that for every set of propositions, there is a proposition which says that every proposition in the set is true. Whatever propositions are, if there are propositions at all then there are sets of them, and for any set of propositions, it is something determinately true or false that all the members of the set are true. If one is willing to talk of propositions at all, one will surely conclude that this determinately true or false thing is a proposition. It may not be possible to express all such propositions since it may not be possible, in any actual language, to refer to all such sets; it may not be humanly possible to grasp them. But if this is so, it is surely a contingent human limitation which should not restrict the range of potential objects of propositional attitudes. So I will assume that the world#story theorist will want to add thesis (W5) to his theory.

Thesis (W6) has some notoriously problematic consequences, but it is also a thesis for which there are independent arguments. [17] If the world#story theorist accepts these arguments, then he will want (W6) in his theory as well. It is compatible with the minimal theory, so he can add it as a further postulate.

(p. 37 ) If (W5) and (W6) are added as postulates to the world#story theory, then it becomes equivalent to the possible worlds analysis with respect to the structure it imposes on the set of propositions. For every set of world#stories, there will be a unique proposition that is a member of just the world#stories in the set. This means that every function from possible worlds (or world#stories) into truth#values will correspond to a unique primitive proposition which is true just where the function takes the value true, and false just where the function takes the value false. It seems that the sole difference that remains between the two theories is that one takes as primitive what the other defines. And even this difference will be eliminated if we make one more change in response to a question about the further development of the world#story theory.

The next question for the world#story theorist is this: can he say more about his fundamental concept, the concept of a proposition? In particular, are there some *basic* propositions out of which all the rest can be constructed? The usual way to answer this kind of question is to model basic propositions on the atomic sentences of a first#order language; propositions are constructed out of individuals and primitive properties and relations in the same way that sentences are constructed out of names and predicates. But this strategy requires building further structure into the theory. There is another way to answer the question which needs no further assumption. We can deduce from what has already been built into the world#story theory that there is a set of propositions of which all propositions are truth#functions: this is the set of strongest contingent propositions—those propositions which are members of just one world#story. It is thus a harmless change, a matter of giving the theory a more economical formulation, to take these to be the basic propositions. (This change does not foreclose the possibility of a further reduction of what are here called basic propositions to propositions even more basic. Any alternative reduction can be expressed as a further reduction of this kind; this is why this move is harmless.) We can then define propositions generally as sets of basic propositions (or, for a neater formulation, call the basic elements *propositional elements* and let their unit sets be the basic propositions. Then all propositions will be things of the same kind.) A nonbasic proposition will be true just in case one of its members is true. This reduction has the added advantage that it allows us to define the previously primitive property of consistency, and to derive all of the postulates. With these primitive notions and assumptions eliminated, the world#story theory looks as good as the theory that takes possible worlds as primitive and defines propositions. This is, of course, because it is exactly the same theory.

I have gone through this exercise of changing the world#story theory into the possible worlds analysis of propositions in order to make the following points. First, the minimal world#story theory with which I began is indeed a minimal theory, a theory that assumes nothing about propositions except (p. 38) what is required to capture the principal intuition of the theory: that something is possible only as part of a possible completely determinate situation. But second, every step in the metamorphosis of this minimal theory into the possible worlds analysis is motivated either by uncontroversial assumptions about propositions (such as that truth#functions of propositions are propositions) or by the pragmatic picture of propositional attitudes, or by theory# neutral considerations of economy. If this is right, then the possible worlds analysis is not just one theory which gives propositions the identity conditions which are motivated by the pragmatic picture. More than this, it is the whole content of that analysis to impose the minimal structure on propositions which is appropriate to a theory which is guided by this conception of propositional attitudes. Anyone who believes that there are objects of propositional attitudes and who accepts the assumptions about the formal

properties of the set of these objects, must accept that there are things which have all the properties that the possible worlds theory attributes to possible worlds, and that propositions can be reduced to these things.

Is the form of realism about possible worlds that I have been defending really realism? It is in the sense that it holds that statements about what is possible are to be explained in terms of quantification over possible worlds, and that some such statements are true. It is in the sense that it claims that the concept of a possible world is a basic concept in a correct account of the way we represent the world in our propositional acts and attitudes. But on the other hand, the moderate realism I want to defend need not take possible worlds to be among the ultimate furniture of the world. Possible worlds are primitive notions of the theory, not because of their ontological status, but because it is useful to theorize at a certain level of abstraction, a level that brings out what is common in a certain range of otherwise diverse activities. The concept of possible worlds that I am defending is not a metaphysical conception, although one application of the notion is to provide a framework for metaphysical theorizing. The concept is a formal or functional notion, like the notion of an *individual* presupposed by the semantics for extensional quantification theory. An individual is not a particular kind of thing; it is a particular role that things of any kind may occupy: the role of subject of predication. To accept the semantics for quantification theory is not to accept any particular metaphysics of individuals, even though one may use the resources of that semantic theory to help clarify one's metaphysical commitments. Similarly, a possible world is not a particular kind of thing or place. The theory leaves the *nature* of possible worlds as open as extensional semantics leaves the nature of individuals. A possible world is what truth is relative to, what people distinguish between in their rational activities. To believe in possible worlds is to believe only that those activities have a certain structure, the structure which possible worlds theory helps to bring out.

(p. 39 ) Let me develop the analogy between domains of possible worlds and domains of individuals to try to make my point clearer. One may think of quantification theory as a framework for representing, once and for all, one's total ontological commitment. To do this is to assume that there is one big domain of discourse—the domain of what there is —which gives the intended interpretation of quantification theory. To be, then, is to be a value of a bound variable, under this interpretation. But one can reject this metaphysical interpretation and still accept the referential semantics for quantification theory. One can deny that there is a domain of all there is, of which all the domains of discourse in particular contexts are subdomains. One can deny that it makes sense to ask ontological questions outside of a particular context.

Analogously, one may choose to put a metaphysical interpretation on the concept of a possible world, assuming that there is one domain of all metaphysically possible

worlds from which the restricted domains relevant to interpreting different kinds of possibility and necessity are drawn. But one may also reject that interpretation, and the coherence of the metaphysical questions which it raises, without rejecting a realistic understanding of possible worlds semantics. One may say that in particular contexts of inquiry, deliberation, and conversation, participants distinguish between alternative possibilities, and that they should do so is definitive of those activities. It does not follow from this that there is a domain from which all participants in inquiry, deliberation, and conversation must take the alternative possibilities that they distinguish between.

Notes:

(1) Mackie ( 1973 : 84).

(2) Powers ( 1976 : 95).

(3) Lewis ( 1973 : 84).

(4) Ibid. 86.

(5) Ibid. 85.

(6) Ibid. 85–6.

(7) Ibid. 85.

(8) Lewis ( 1979 : 533).

(9) Augustine (399/ 1943 : 276); (Bk. 11, ch. 20).

(10) Lewis ( 1973 : 88).

(11) Mackie ( 1973 : 90, 92).

(12) Adams ( 1974 : 224).

(13) Ibid 225.

(14) Adams ( 1974 : 224).

(15) In the earlier published version of this paper I claimed otherwise. Philip Bricker, in his dissertation (Bricker 1983 ), pointed out my mistake in assuming that (W4) followed from (W1)–(W3). Bricker's dissertation contains a clear, detailed, and interesting investigation, both technical and philosophical, of the relation between proposition#

based theories such as Adams's world#story theory and world#based theories such as the analysis I am defending.

(16) Adams ( 1974 : 228).

(17) Both the problematic consequences and the independent arguments are discussed in Stalnaker ( 1984 : ch. 1).