A Fundamental Theory of Genuine Modal Realism*

© 2005 by Damon Woolsey

I

It is difficult, if not impossible, to construct an argument in favor of Genuine Modal Realism (GMR), the thesis that all possible worlds are just as real as our own. GMR seems to be the utmost in metaphysical extravagance, postulating an infinite number of other worlds of which we can have no direct experience. Common sense balks at such a proposition, and for the very same reason, one could invoke Occam's razor, since these other worlds seem to play no causal role in our own world. Nevertheless, I am a firm believer in them; *all of them.*

I regard the principle of simplicity as something of an a priori constraint on theorizing, and therefore tend toward strict adherence to it. The problem here is that GMR will always seem, *prima facie,* to multiply entities beyond necessity. The qualifier here is necessity, and so we must consider what is in fact necessary. Most regard our world as contingent, in which case it could have been different from how it actually is, or perhaps may not have existed at all. If our world truly is contingent, then why does it exist? To say that it could have been different is to say that there is some possible world that *is* different, and that it could have been the actual world. So what is this property of "being the actual world"? Such a question seems to presuppose the existence of only one world, the actual world, and so we should rephrase the question in terms of being a real world. Certainly, any possible world could be a real world. We may then wonder how it is that our world has achieved the prominent status of "being real" while other possible worlds have not. There is a sense that for a possible world to be real like our own, it must be built from "real stuff". What this stuff is and where it comes from is currently unknown, but I am

fairly certain that it isn't manufactured in some sort of cosmic factory. Perhaps, if we are to level a charge of metaphysical extravagance, it should be against the notion of "real stuff", rather than the reality of all possible worlds.

There are two things that could be meant when we say that our world may not have existed at all. In the literal sense, it means that our world would not exist. This would seem to reduce the status of our world to that of other possible worlds, for even if our world was not real, it would still be a possible world. We could also mean that our world may have been an empty world wherein nothing exists. But that is not our world, for our world, whether it were real or merely possible, contains exactly what it does. Perhaps there is some other possible world that is empty (cf. Rodriquez-Pereyra), but not ours. So if our world did not exist as a "real" world, then it would still exist as a possible world. Again, we are back to the question of how a possible world becomes a real world. I don't think that there is any difference between being possible and being real, and it seems to me that the burden of proof is on anyone who thinks that there is. What I am denying here is the distinction between worlds that are somehow "real" vs. worlds that are "merely possible", which is quite different from the distinction between the *actual* world and other possible worlds. The term "actual" is an indexical, and simply refers to the world of which one is a part.

David Lewis was by far the most famous and outspoken defender of GMR. His book, *On the Plurality of Worlds,* served as an explication and defense of his version of GMR, which I refer to as *Lewisian* GMR. His efforts did little to stem the tide of objections, however. It seems that GMR, as developed by Lewis himself, is indefensible. But here I intend to offer a reformulated version of GMR that I believe to be far superior to Lewis'. Rather than rebutting the vast array of objections, on this new version they simply carry no weight. Furthermore, by doing

away with certain unfounded and arbitrary suppositions, a coherent account of possible worlds emerges.

The intuitive concept of a world is much as Lewis described it: a world is a concrete mereological sum of its parts. Everything within a world is spatio-temporally related to everything else, and every part of it is contained wholly within it. As such, worlds are isolated, since no two worlds share any parts in common, and there can be no trans-world causation. Thus, one of the major tenets of Lewisian GMR is that worlds do not overlap. This supposition has important consequences for his characterization of GMR.

There are two principle ways of knowing about other worlds: *counterpart theory* and the *principle of recombination.* In some other worlds, there are individuals very much like yourself. These individuals are your counterparts. For any event in your life that could have been different, there is a counterpart in some other world for whom it *was* different. This extends to others in our world as well; we could, for example, ask what would have happened if Hitler had been accepted into art school. There is some other possible world wherein Hitler's counterpart was accepted into art school, and history turned out very differently. The other method of knowing about other worlds, the principle of recombination, simply allows for the unrestricted recombination of parts of worlds. For example, since our world contains horns and horses, there is some possible world wherein there is a unicorn.[1] Or, since some other world contains hobbits, and our world contains humans, there must be some world that contains both humans and hobbits.

So much for my all-too-cursory synopsis of Lewisian GMR. I needn't go into too much detail here, but shall proceed by critically examining the major tenets of Lewisian GMR and revising them as needed. These revisions were not entirely foreign to Lewis; as we shall see, he

was open to considering them all. His rationale for rejecting them seems to have been based on common sense, but common sense is fallible, and should be eschewed if there is good enough reason to do so.

<center>II</center>

That worlds are isolated follows logically from the very definition of a world, for if some individual a were part of both worlds w_1 and w_2, then by transitivity of relations, every part of w_1 would be spatio-temporally related to every part of w_2. But this definition of a world, and the ensuing requirement of isolation, leads to an excess of worlds. The problem is that, at any given point in time, there are an untold number of worlds that are identical to our own. These worlds *will* differ from ours at some point in the future, but until such time they follow the same course as our world.[2] Likewise, there are a great number of worlds that *were* identical to our own, up to some point in history, whereupon they diverged. But why should we accept an infinite number of identical copies of our world, just so they can differ from ours at some point in time? Perhaps such a requirement is a relic of the notion that worlds need to be manufactured out of "real stuff". It seems simpler that, if two worlds are identical at some point in time, they should be the same world at that time. After all, how can two worlds be different if there is no difference between them?

Lewis was not entirely incurious toward overlap of worlds: "… limited overlap might be wanted. The most likely case would be limited overlap when branching worlds share a common initial segment" (1986, 206). This is the sort of overlap that I propose, but it requires redefining the term "world". I prefer to think of a world as something like a game: it begins with a number of pieces (i.e. fundamental parts) and a set of rules (i.e. natural laws).[3] There are an infinite

number of different types of games that can be played according to the various combinations of pieces and rules. Each type of game corresponds to a *tree structure*, and together they (i.e. all possible worlds) form a *forest*.[4] The root of each tree is the initial state of the game. During any given turn (time) there are a number of different possible states of affairs for any game. What distinguish branching worlds within the same tree are the different possible states of affairs between them.[5]

Here we must abandon the notion that temporal relations are restricted to a single world, for one initial world branches into many, and they are all temporally related to the initial state. But spatial relations, by themselves, are not sufficient to determine a world, because then a world would consist entirely of a single time-slice of something bigger. However, from our tree structure we can trace a world line by connecting a unique series of consecutive states of affairs, and that is sufficient for my purpose. There is a sense in which, according to my theory, a single spatio-temporally related world is just a fiction. But the common sense notion of a world is something that needs to be explained, and the discussion that follows will help to clarify this matter.

This account of branching worlds has many benefits, but Lewis had reason to reject it: "Branching, and the limited overlap it requires, are to be rejected as making nonsense of the way we take ourselves to be related to our futures; and divergence without overlap is to be preferred." (1986, 208) It seems most natural that, although most regard the future as indeterminate, we have but one future before us. Branching worlds entails that not one, but all possible futures are equally ours. This does indeed make nonsense of such questions as, "Where will I be five years from now?" Perhaps such questions really are nonsense. However, while Lewis offers an intuitive argument against branching, there is an equally good argument for accepting it. We do

indeed regard the future as indeterminate. But if worlds are isolated, and all possible worlds exist independently from the beginning, then your future is strictly determined by the world of which you are a part. Different futures belong not to you, but to your counterparts, and there is but one course your life can take. This course was set out before you were born, and there is nothing you can do to alter it. I think most would reject such determinism in favor of indeterminacy, but the problem of which future is yours remains.

For some this dilemma is reason enough to reject GMR altogether. It does indeed baffle the mind to think that all possible futures are equally yours, but I contend that this sort of GMR actually makes sense of the way we think of our futures. When I think of my past self, I think of but one past, the one that actually happened to *me*. However, when I think of my future, it is pure speculation. I imagine myself in various places, and doing various different things. I may prefer some such fantasies to others, but I have no right to claim a single one as being "mine", while discarding the rest as belonging to my counterparts. The real difficulty here is that we are dealing with identity over time, and whenever we examine the notion of personal identity we encounter trouble. This is not strictly a problem for GMR, but it does add insult to injury. It would help if we had a better notion of what "I" means within the framework of branching worlds.

When one uses the term "I" to refer to oneself, it is always with a temporal qualifier. For example: I was, I am, I will be; I did, I am doing, I will do; or, I thought, I think, I will think. When making reference to oneself, the benchmark always seems to be here and now. When I speak of my past self, in so far as my memory serves me, it is always from the reference point of the present. Sometimes I may be mistaken about my past self, or not remember certain things at all, but this causes no concern for my own personal identity here and now. But what about when I reference my future self? What exactly am I making reference to? To put it another way: what

sort of epistemic access do I have to an indeterminate future me? If Lewisian GMR is true, then the future course of our world is very strictly determined, and there is a definite future "I".[6] This conflicts with the modern scientific view that the universe is in fact indeterminate, in which case there is no determinate future "I" to refer to. On the other hand, if GMR with branching worlds is true, then there is no single future "I". This preserves our notion of an indeterminate future, but is still not quite in line with our intuitions. It will help to discuss what it means for the "I" of the present to become the "I" of the future.

One's sense of identity is always rooted in the present, and at any given time there is only one "I". As events unfold around us, we find ourselves at a point that we once regarded as the future. But this once-future-now-present is still exactly where we find this elusive "I". The present is entirely inescapable; no matter how long you wait, it is still *now*. So the real objection isn't about who we are, it is about who we want to be. Some (cf. Beedle) argue that GMR entails fatalism. Why should we look both ways before we cross the street, when we're just going to get hit by a car in some possible world anyway? The answer is quite simple, and follows from my definition of a world. Events unfold according to the causal rules of our world, and as an agent you have a degree of control over what is in fact a possible future for you and what is not. Indeed, is that not how we plan for the future? We make plans and do our very best to carry them to fruition, but all the while we are preparing for contingencies. I tend to think that a cautious person, when crossing a busy street, fares better in all possible worlds than a reckless buffoon.

III

This account of branching worlds still cannot quite explain what a "world" is. Lewis was aware of this difficulty: "Thus overlap complicates what we must say in explaining how worlds

are unified… The complication is unwelcome, but I think it's nothing worse" (1986, 209). Thus far I have been casual in my use of the term "world", and shall continue, for now, to use it with the common sense meaning, but I also mentioned that a single spatio-temporally unified world is something of a fictional entity. Objectively, there is no such thing, but subjectively, from the point of view (POV) of an observer, there certainly seems to be. What is needed is an account of how an objective states of affairs view relates to a subjective view of a world. In other words, we need to sort out our philosophy of mind from the rest of our metaphysics. This complication arises from the way we normally think about our world. In the end, however, when all the difficulties have been ironed out, what remains is actually quite simple and elegant.

When I mention the POV of an observer, it is not my intention to propose some special metaphysical role for conscious minds. And while conscious minds are an important topic in metaphysics, I do not intend to offer a theory of mind here; that is not the subject of this paper. Nor do I intend to discuss the matter of consciousness. In the way of sorting things out, however, I will have something to say about these two subjects. Consciousness, I assume, supervenes on certain facts about our world. Whether these are physical facts, or further facts about our world, matters not to this discussion. What is pertinent is that consciousness is a product of the natural laws of our world, rather than some extra-worldly power (e.g. some deity). As for minds, how they work is not something I need go into here. It will suffice to say that a mind is instantiated by a causal process at work within a world.[7] As a causal process, it follows the same natural laws as everything else. So what is really pertinent here is not conscious minds, but causal processes.

Let us call the causal interaction of two or more things a *nexus.* Such interactions can include a photon slamming into an electron, an asteroid circling the sun, or the complex interaction of a neural network. All of these create nexuses; the last sort creates a particular sort

of nexus known as a POV. The specific states of affairs within a single nexus are not sufficient to determine the states of affairs of an entire world, for a single photon from our sun colliding with an electron in the earth's atmosphere has nothing at all to do with the specific states of affairs in a far-off galaxy.[8] Or, my sitting here typing this paper is completely unrelated to the location of a particular grain of sand in the Gobi desert. This is not to rule out the Butterfly Effect, by which small factors can have profound influences farther down the causal chain of events. What I speak of is a *simultaneous* relation, wherein the states of affairs of two nexuses have no *immediate* causal influence over one another. Other intermediary nexuses buffer them from one another, thereby yielding multiple compatible possibilities. If, however, I traveled to the Gobi desert, and a particular grain of sand were blown into my right eye, that *would* have an immediate effect upon me.

Now, if we were to imagine two possible worlds, having differing states of affairs between them, there still may be certain nexuses within them that have identical states of affairs. For example, there may be one world wherein a particular grain of sand is on the west side of the Gobi desert, and yet another where it is on the east side. However, in both worlds I am doing exactly the same thing: typing this paper. I propose, yet again, that such duplication is unnecessary. We can imagine these two worlds as though they exist independently, but really there is just one of me, typing this paper. The fact that there are innumerable possible locations for that grain of sand should not mean that there must be innumerable identical copies of myself. It makes no immediate difference to me where that grain of sand is, and so from my POV, the grain of sand is a mere *possibilia*. Likewise, it makes no difference to the grain of sand what I am doing at the present time. Any two nexuses may coexist in a state of mutual possibilia, so long as the various possible states of affairs in one are logically consistent with the possible

states of affairs in the other. This condition of being logically consistent needs more clarification, for there are two factors involved: their immediate coexistence, and their etiological coexistence. The first is an account of overlapping possibilia; the second serves as a limiting factor on them.

The immediate coexistence of several objects with one another is a fairly straightforward matter. Imagine that I have three marbles: one each of red, blue, and green. I have placed one marble in each of three urns labeled *x, y,* and *z.* There are exactly 3! (six) possible combinations of marbles and urns (provided that there is exactly one marble in each urn). Since you have no knowledge of which marble is in which urn (i.e. they have no immediate causal influence on you), each of these six possibilia is mutually consistent with your POV. Suppose that you then look into urn *x* and see that it contains the blue marble. Now there are only 2! (two) possible combinations, for it would not be logically consistent for the blue marble to also be in one of the other two urns. Upon looking inside urn *y* you then see that it contains the red marble. There is now only one possible state of affairs for the three urns that is logically consistent with your POV, and so the green marble must be in urn *z.* Before you looked in any of the urns, each of the six possibilia was logically consistent with your POV. It would be senseless to claim that, because there were six possibilities, there were six of you, since every duplicate of yourself would be exactly identical. Upon looking in the first urn, however, the contents thereof had an immediate causal impact on you, leading to three different possible states of affairs of yourself: one where the marble was red, one blue, and one green. For each of these three possibilities there was one of you (for a total of three of you), and for each one of you the remaining possibilia for the total state of affairs of the urns was two.[9] Lewis would regard this as overlap of worlds, but as I have argued, spatio-temporally unified worlds are mere fictions. Worlds don't overlap, but possibilia do, and this overlap occurs from nexuses. The number of different possible states of

affairs, in relation to an individual nexus, is given by the permutation of the individual possibilities. From a nexus, possibilia must be logically consistent, and immediate coexistence (as when you look in an urn, or a photon collides with an electron) limits the number of logically consistent possibilia.

Immediate coexistence, over the course of events, serves as a constraint on what is in fact possible. Suppose that we repeated the above exercise, but this time I (truthfully) told you that, regardless of which marble I placed in the first urn, I would place the rest in alphabetical order by color. For example, the order may be blue-green-red, green-red-blue, or red-blue-green. This prior condition serves as an etiological constraint on the possibilia and limits them from six to three. But that is not the end of it, for upon looking into the first urn, this prior constraint remains in effect, limiting the remaining possible permutations to one. So long as such a prior condition is relevant, it will continue to serve as an etiological constraint on possibilia. It is important to note that such etiological constraints are not themselves causal influences; they play no part in causation.[10] Etiological and mutual coexistence are strictly logical relations between mutually consistent states of affairs.

IV

This account of overlapping possibilia from mutual coexistence with a given nexus is far removed from the clean tree structure of branching worlds. What we now have is more like a web of mutually consistent nexuses. There are a number of different possible states of affairs within any given nexus. For each possible nexus, there are a number of mutually coexistent possible states of affairs of other nexuses. A great number of these overlap, but due to etiological constraints some of them are not mutually consistent. In this tangled web it is quite difficult to

determine exactly what a world might be. Branching worlds disturbed our notion of temporal

isolation, and now overlapping possibilia have obliterated spatial isolation. Still, there remains a

simple way to define a world that keeps with our intuition. A world is simply a *maximally*

consistent (i.e. both mutually and etiologically consistent) state of affairs over time.

Theoretically, we can achieve such a conceptualization by choosing one possible state of affairs

for each nexus, provided that they are all logically consistent with each other.[11]

Now that we have a clear idea of what a world may be, we can now turn our attention to

knowing about possible worlds. Lewis gave two primary methods for knowing about possible

worlds: counterpart theory and the principle of recombination. The former is an account of how

an individual can be, while the latter is an account of how a world can be. Of course, this does

not provide a method for determining how all the worlds are, for surely there are worlds with

unfamiliar things obeying unfamiliar natural laws that cannot be understood in terms of what

exists in our world. These unfamiliar worlds are called *alien worlds,* and if GMR is true, then

there is an abundance of them. Accounting for them is a serious difficulty that I needn't go into

here. Rather, I shall attend to the topic of *counterfactuals* in relation to the things and events that

we are familiar with. This is the practical application of GMR, whereby we can speak of how

things could have turned out differently.

Lewis' counterpart relation works according to qualitative similarity in content and

context, where whatever thing in one world is most similar to a thing in another world is its

counterpart (1968, 114). However, it seems that such a similarity relation is far too weak to

account for full-blown *de re* representation, since similarity admits of degrees; it is not a set of

necessary and sufficient conditions. There is no non-arbitrary method for establishing what

would be similar enough to be a counterpart of something and what would not.[12] It would seem

that there must always be some thing in any given world that bears some sort of similarity to myself, and so there must always be some thing in that world that is most similar to me (except, perhaps, in the case of alien worlds). How then could we say that there is a world wherein I have no counterpart? Say, for example, a world in which my parents never met. We would want to say that I don't exist in such a world, but by a similarity counterpart relation, I could still seem to have a counterpart. Lewis states that if nothing in a particular world resembles me closely enough, then I have no counterpart in that world; but what counts as being close enough? (1968, 116) *De re* representation should be non-arbitrary, and in any event qualitative similarity sounds more like *de dicto* representation than *de re* representation.

The situation is even worse for the principle of recombination. Lewis claims that for any combination of parts of this world, there is some world that is that way. He evens goes so far as to discuss a world containing disembodied talking heads! (1986, 87-88) But there is a more serious difficulty for such unrestricted recombination: there must be more worlds than there actually are. This is a paradox of cardinality, and I won't repeat the argument here (cf. Lewis, 1986, 102). The important thing is that Lewis' solution to this problem is that there must be some limit to the possible size of spacetime, which seems entirely *ad hoc.* Lewis doesn't claim to know what this limit is, but he holds out the hope that there is some natural mathematical break that will be obvious when it is discovered (1986, 103). Such a natural break would not be arbitrary, but it doesn't exist. Mathematics generalizes to fit any arbitrarily large system. There is no largest integer, nor a limit to the decimal expansion of the real numbers, and geometry is extensible to any arbitrary number of dimensions. There is no object so large or existing in so many dimensions that it cannot be described mathematically. Lewis' restriction on the possible size and shape of spacetime is therefore to be dismissed.

The solutions to these problems with counterpart theory and the principle of recombination follow naturally from the account of overlapping possibilia from mutual coexistence with a given nexus. As events unfold, what is mutually consistent is etiologically constrained by the different possible states of affairs of the nexuses involved. This was illustrated by the example of the marbles and urns. Another example will make clear how this relates to counterpart theory. Suppose that I am sitting here typing this paper, and my doorbell rings. Certainly, there are a vast number of different possibilities regarding who is at my door. Each one of these possibilia is consistent with my POV, so they are overlapped, but there is only one of me sitting here. However, from the POV's of the possible individuals ringing my doorbell, they are not overlapped, for that would not be mutually consistent between them. Each one of them is standing there alone, ringing my doorbell, and I am here listening to my doorbell ring. The doorbell serves as an intermediate nexus, thereby buffering me from them. When I open my door, however, it would not be consistent to see them all standing there, overlapped. These are different states of affairs that are not mutually consistent. At this point of immediate interaction between overlapping nexuses, there must be one of me to greet each of the different people who could be ringing my doorbell. This is how nexuses branch; it is a strictly logical relation.

This branching fixes the counterpart relation in a way that similarity cannot. When considering a counterfactual situation, we don't have to "look" in another possible world for an individual that bears the greatest similarity to an individual in our world. All that is needed is to consider the individual who is part of *our* world, and examine the point where these worlds branched. This involves identifying an event that could have turned out differently than it did in our world, which we do by stipulation. Then we simply ask what would have happened to the individuals involved had that event turned out differently.[13] There is no question here as to which

individual in that other world is the counterpart of a given individual in our world; it is simply

that individual in a counterfactual situation. This is in agreement with Kripke's claim that "things

aren't 'found out' about a counterfactual situation, they are stipulated", and "possible worlds need

not be given purely qualitatively, as if we were looking at them through a telescope" (pp. 49-50).

Furthermore, "We can point to the *man*, and ask what might have happened to *him*, had events

been different" (author's emphasis) (p. 46). In this last case, Kripke is using the past tense to

describe an event, and that is what we have to do: look to the past.[14] From there, we can identify

an event that could have turned out differently, and ask what events would have followed.

The view that I have laid out has important consequences for the principle of

recombination, because, due to the etiological constraint, not every combination of things is

possible. For a world to contain a given combination of things, there must be some chain of

events that led to that specific combination. And if that combination of things involves a

counterpart of some thing in our world, then there must be some preceding event(s) that could

have turned out such that *that* combination came to be. If there is no such chain of events, then

there is no such counterfactual situation involving a counterpart of something in our world. Thus,

when we use the principle of recombination, it must be etiologically constrained by events. I call

this the *causally constrained principle of recombination*. This also limits the general case of

recombination (e.g. where some things that are like horses also have horns on their heads, which

we could call unicorns). Here there would need to be some world where evolution worked out

differently, and I think that is possible. But this different evolutionary history would have other

consequences beyond the existence of unicorns. Perhaps, in such a world, humans would not

have evolved because things would have been *too* different. Or, maybe if humans had evolved,

you wouldn't have been born. So, if you imagine a possible world wherein you are riding a

unicorn, you may actually be imagining an impossible world. A more obvious constraint is that it

is impossible for a state of affairs to exist that is self-contradictory. There is simply no chain of

events that can lead to a contradictory state of affairs. One example of this is Lewis' talking

heads. How can a head talk if it doesn't have lungs? It can't! Our ability to imagine impossible

worlds is quite remarkable, but we shouldn't confuse imagining with possibility.[15]

Possible worlds that did not branch from our world cannot account for such impossible

recombinations either. They are alien worlds, and do not contain any counterparts to anything in

our world. To use my game analogy, they have different pieces and/or operate under different

rules. These worlds aren't anything like ours, and certainly do not contain anything that we

would be familiar with (like horses, horns, and unicorns). There is little, if anything, that can be

known about them. John Divers has argued that the existence of alien worlds constitutes an

analytic limit of GMR, and I agree (cf. Divers). The only counterfactual situations that we can

consider involve those worlds that branched from ours, and these are logically constrained, both

immediately and etiologically.

V

I think that modal realism is one of the most abused tools of philosophers, who have a

tendency to use it for backwards reasoning. Take, for example, the notions of "possible" and

"necessary". The typical modal definitions of these terms are that something is possible iff it

exists in some possible world, and something is necessary iff it exists in all possible worlds.[16] It

would be nice to be able to account for possibility and necessity in non-modal terms, and I think

it is the case that something that exists in some possible world does so *because* it is possible.[17]

Likewise, a necessary proposition is true in all possible worlds *because* it is necessary. This is

how worlds are formed. Of course, we can infer possibility and necessity from the existence (or truth) of an individual (or proposition) in some or all worlds. However, GMR simply cannot say what is possible and what is necessary; that is an entirely separate metaphysical question. This may limit the practical uses of GMR, but it is nothing we need worry about. Modal realism, ersatz or genuine, has always had this difficulty. When has anyone ever verified something's necessity by looking in every possible world to see if that thing exists? That would be pure nonsense. It has always been the case that we assumed some non-modal basis for possibility and necessity. I don't intend to delve too far into these issues here, but I do think that 'possible' simply means that it is not self-contradictory, and 'necessary' means that its negation implies a contradiction.

This version of GMR, then, does not purport to offer some easy analysis of possibility and necessity; nor does it attempt to provide the content of all possible worlds through recombination, due to the existence of alien worlds. Possibility is a matter of non-contradiction, and this isn't always so obvious. Certain things, such as square circles, are obviously self-contradictory, and hence impossible; we can't even imagine such an object. However, it is often the case that the inherent contradiction is hidden away in the complexity of the object being imagined. Recall Lewis' talking heads. We can imagine such a thing, propped up on a desk perhaps, talking to us about this or that. But when we imagine it talking, it is talking in the normal method of expelling air from lungs it does not have. This is a contradiction, and hence such a thing is impossible. What would be possible, however, is for there to be some mechanical device that looked like a disembodied head, which could emit sounds through a speaker in the back of its mouth. I am not certain whether some such mechanical device could be replicated biologically, and am much less certain that such a thing could evolve naturally. Perhaps some

mad biologist could produce such a talking head in a laboratory, but that seems unlikely. The point of all this is that the possibility of an object is not something that need be obvious to us, for we can easily imagine impossible objects as well. Therefore, it is not the case that everything we can imagine must exist in some possible world. Likewise, the principle of recombination is not a reliable guide as to what may exist in other possible worlds, for it provides an incomplete picture. Is there some possible world that consists entirely of an oak tree? I think not, for a tree requires nutrients, soil, an atmosphere, moisture, and sunlight; it cannot exist singularly. The ability to imagine something is a poor indication of its possibility, and the reckless use of modal operators makes for unsound arguments. So not only is there an analytic limit to GMR, but there is an epistemic limit as well. This epistemic limit means that determining whether or not an imagined state of affairs is a possible state of affairs can be a tedious undertaking. This is acceptable because GMR really isn't a tool for determining what exists; it's just a model of how reality works. There is no practical benefit to accepting the validity of GMR, except as a structural framework for understanding various observations we make about our world.

VI

I find it interesting that both the many-worlds interpretation of quantum mechanics and genuine modal realism postulate a plurality of worlds, and given that they both purport to describe reality, I think they must both be describing the same thing (cf. Kripke, 45n). While physicists and philosophers may be said to live in their own little worlds, it would be quite strange if there really were one reality for physicists, and yet another for philosophers. If there is indeed but one reality of which we all share (and of course there is), then it seems reasonable that a philosopher's account and a physicist's account should be reconcilable with one another. I will

simply state, without elaboration, that I believe this version of GMR to be compatible with quantum physics, and that it is all the better for it. It is difficult to form a robust sense of reality in the face of quantum weirdness that allows us to keep our common sense intuitions, but here I presume to have done not only that, but one better. This version of GMR explains the greatest quantum mystery of all: entanglement.

Two particles can be correlated such that over arbitrarily large distances a measurement of some property of one induces an immediate corresponding state (which had been previously indeterminate) in the other. Such an instantaneous connection seems to violate the rule that nothing can travel faster than the speed of light. The question is, how does one particle instantaneously "know" what the other particle is doing? Einstein referred to this as "spooky action at a distance", and it was one of his primary reasons for claiming that quantum physics is incomplete. According to general relativity, no particle or force can travel faster than light. How then can two particles enjoy such instantaneous correlation? Quantum mechanics offers no explanation for this phenomenon, and physicists have searched in vain for some physical cause of entanglement. I predict that there is no physical cause, because entanglement is not a physical relation; it is strictly a logical relation, and is due to the etiological constraint that I outlined in part IV. If two particles were in the same state, even if that state is indeterminate (i.e. overlapped possibilia), then it would not be consistent to find one in one state and the other in a different state. This etiological constraint would only be observable if GMR with overlapping nexuses is true, and as it is capable of explaining quantum entanglement, an experimentally verifiable phenomenon, I take it as evidence that it *is* true.

Another quantum mystery is the nature of indeterminacy. It seems that no event is strictly determined, since certain conjugate properties of particles cannot be precisely measured (e.g.

position and momentum). So a quantum system gives rise not to just one possible state of affairs, but a multitude of them. The problem is that, if ours is the only "real" world, then what happens to all the possibilities that did not manifest themselves? How is it that, out of all the possible states of affairs, only one becomes real? Is the universe truly that arbitrary? GMR doesn't have to answer such questions, for it suffers no arbitrariness; all possibilities happen. Such a view restores strict determinism on a grand scale, and in a way that keeps with the intuition that our world is indeterminate. No event is random on this view, and yet from within any given world, "random" events do happen.

Another difficulty that is resolved is that of an uncaused first cause. The initial state of a world is just some number of parts and a set of natural laws to abide by. This version of GMR says nothing about initial position and momentum, which might seem quite strange. However, is that not how our world began? According to the big bang theory, all the matter in the universe was condensed into a point of infinite density. The resultant explosion was a direct result of that initial state. So our question is, what caused that initial state? The common notion of an uncaused first cause is something that existed prior to our universe, which then caused our universe to come into existence. Theists take it to be God, while non-theists either ignore the problem or state that such a question is beyond the realm of scientific inquiry, since physics breaks down at times less than 10^{-43} seconds after the big bang. However, while this question may be beyond physics, it is not beyond *meta*physics. While most assume that an uncaused first cause was the act of creating or the event of becoming of our world, I claim that the uncaused first cause was simply the initial state of our world. This may seem unsatisfactory, since you wonder, "Where did all that stuff come from in the first place?" Again, we are back to the question of how a possible world becomes a "real" world, and the intuition that real worlds must be built from "real

stuff". I deny this intuition; possible worlds don't need to be built out of anything at all, and so there is no difference between some particular thing being possible, and it being real. The initial state of a world simply is, by virtue of being possible (i.e. having no inherent contradiction). What this amounts to is that possible worlds necessarily exist.

I don't presume to have convinced anyone of any of this, except myself. There is always that incredulous stare, and the corresponding outright dismissal of GMR as ludicrous. I don't consider that to be a valid objection; however, given such an outrageous claim as GMR, it doesn't really have to be. I can see no reason for accepting GMR as being true, except for one: it makes perfect sense.

References

Beedle, Andrew; Modal Fatalism; *The Philosophical Quarterly,* Volume 46, Number 185, pp. 488-495; October 1996.

Divers, John, Joseph Melia; The Analytic Limit of Genuine Modal Realism; *Mind*, Volume 111, Number 441, pp. 15-36; January 2002.

Kripke, Saul; *Naming and Necessity;* Basil Blackwell Publisher, Oxford; 1980.

Lewis, David; Counterpart Theory and Quantified Modal Logic; *The Journal of Philosophy,* Vol. 65, No. 5, pp. 113-126; March 1968.

— *On the Plurality of Worlds;* Blackwell Publishing, Massachusetts; 1986.

Rodriquez-Pereyra, Gonzalo; Modal Realism and Metaphysical Nihilism; *Mind*, Vol. 113, No. 452, pp. 683-695; October 2004.

Notes

1. This example is not meant to contradict Kripke's argument that these things wouldn't be unicorns. If you like, we could call them "things that look like unicorns" or "things that look like horses with horns".

2. On Lewisian GMR, there is some difficulty in establishing temporal correlations between worlds, since worlds are spatio-*temporally* isolated. I pass lightly over this difficulty here because it will soon be resolved. However, we can (as Lewis does) establish some such relation of times between worlds due to the identical states of affairs between them (cf. Lewis, 1986, 70-77).

3. Here I am purposefully oversimplifying the metaphysics. This version of GMR is consistent with a variety of views, including physicalism, neutral monism, and property dualism.

4. Such tree structures, and the notion of a forest, are the same sort as used in the field of computer science.

5. This treatment of branching is grossly oversimplified here, but will be laid out in more detail later.

6. Lewis accepted a sort of epistemic indeterminacy, which is wholly inadequate in light of quantum physics. It may be possible to accommodate metaphysical indeterminacy from within Lewisian GMR, but to do so is to accept indeterminacy of the future, which was Lewis' primary reason for rejecting branching worlds. In other words, if you want metaphysical indeterminacy, branching worlds is the way to go.

7. I don't intend to go into what causality is either. I simply assume that there is some such thing.

8. We may say that a nexus under-specifies a world. Thanks to Xxxxxxx Xxxxx for suggesting this phrasing.

9. Thus, on this version of branching GMR, possibilities are conserved. They don't magically disappear, or collapse into one "real" state of affairs. Call this the Modal Law of Conservation of Possibilities.

10. The etiological constraint is due, roughly, to the causal history of a chain of events. For some imagined or stipulated state of affairs to be possible, there must be some chain of events that would have lead to that particular state of affairs. Using the example from above, if I put the marbles in alphabetical order, there is no possible world where their order is green-blue-red. The importance of this etiological constraint will become apparent in part VI.

11. To achieve a maximally consistent state of affairs would amount to a Laplacean view of the universe, where the properties of each particle were precisely known. Quantum physics says that this is impossible.

12. Lewis admits that his counterpart relation, being based on similarity, suffers from all the usual problems thereof (1968, 115).

13. Here we must be careful that what we are stipulating is in fact possible.

14. It is something of an added benefit that, on my view, temporal relations across branching worlds are not a problem. Temporal relations are easily established by branching, since two branching nexuses were identical at some time. This is not possible with Lewisian GMR.

15. Lewis was vaguely aware of a broader constraint on recombination, but didn't have any proper way to formalize it within his counterpart theory (cf. 1986, 88-89).

16. This extends to propositions also, where a proposition is possible iff it is true in some possible world; and a proposition is necessary iff it is true in all possible worlds.

17. I don't think that any particular individual is necessary, however, for there should always be some world where that individual doesn't exist at all. Certain propositions, on the other hand, are necessary.