

A Genuine Modal Realist's Solution to the Problem of Advanced Modalizing

1. Introduction.

Cases of so-called 'advanced modalizing' are problematic for genuine modal realists in two big ways: they call into question the adequacy of the standard Lewisian translation schema for modal sentences, and (perhaps worse) they indicate that genuine modal realism fails as an analysis of modality. So far, those seeking to defend genuine modal realism have chosen either to revise the standard translation schema, or to recast genuine modal realism as a *non*-reductive account of modality. In this paper, I suggest a solution that leaves genuine modal realism, its translation schema, and its ontology intact. I start with a refresher of the basics of Lewis' modal realism (section 2) that will allow me to present the problem of advanced modalizing clearly and succinctly (section 3). I will then offer give my solution to the problem in broad brush strokes, before turning to look at some of the intricacies of my account (section 4). Where notation is used, I will follow Lewis 1968: $Wx = x$ is a world, $BSx = x$ is a blue swan, $Ixy = x$ is in possible world y .

2. Background.

Advanced modalizing is a problem for genuine modal realism. That is, it is exclusive to the modal realism of David Lewis. That said, you don't have to agree with everything that Lewis says in order to fall prey to the problem. There are various aspects of his view that aren't necessary for generating the problem, such as his specific view about sets or properties. So I'll start by specifying the kind of view that we are considering when we are grappling with the problem of advanced modalizing.

Nowadays it's very common to talk about what's possible and what's necessary in terms of possible worlds: a proposition P is possibly true if and only if it is true at some possible world, necessary if and only if it is true at all possible worlds, and impossible if P

fails to be true at any possible world. This way of talking is a very convenient analytic tool that helps us to sort out modal claims in various fields and its usefulness is acknowledged in almost every subfield of philosophy. We all know that possible worlds talk is a powerful tool. A big question, however, is how exactly we should understand our talk of possible worlds. After all, what *is* a possible world? Their job is to represent the ways things might have been, but when we talk about worlds do we really want to reify something so mysterious as ‘ways’?

Part of the project in the modality literature has been to answer the question of exactly what we are talking about when we talk about possible worlds. Lewis’ modal realism answers that question in perhaps the most radical way possible. He tells us that the possible worlds we talk about are not representations of things or stories that we might tell, they are real, concrete, universes that exist in just the way our universe exists. That brings us to the first important component of the Lewisian view we will be considering, its *ontology*.

Lewisian modal realism (and I will use the terms ‘genuine modal realism’, ‘Lewisian modal realism’ and just ‘modal realism’ interchangeably from here on) claims that possible worlds exist, and (and this is crucial) that those worlds *are ontologically on a par with the actual world*. The other possible worlds are ‘concrete’ (in the sense that they are, ontologically, just like our world), and ‘spatiotemporally isolated’ (in the sense that there are no spatial or temporal relations between things in separate worlds). Each individual world is its own complete universe, each representing a different way that things might have played out. None of the worlds can interact with one another, and no two worlds are fused or combined in any way.

It’s important to note that this is not the only way to answer the ‘what are possible worlds?’ question. Other views, such as linguistic ersatzism claim that possible worlds are *linguistic representations* of the way the world might be: descriptions of worlds in which things happen differently. On such a view, the actual world is ontologically very different

from the other possible worlds, after all it is real and they are abstract. For Genuine Modal Realism, possible worlds are ontologically robust in an important sense: they truly play out the different ways that our world might have been. For any thing that might have happened at the actual world, there is a world where that thing really does happen – not a world that describes or represents that event as happening, a world where that event really does happen.

The first component to our Lewisian view is the ontology. A second important component that needs to be acknowledged is its *reductive character*. The ontology of Lewis is a hard to swallow, but the argument is that it is worth swallowing because it allows Lewis' realism to be reductive in nature. Because the possible worlds are concrete and nonmodal (as opposed to descriptions that are irreducibly modal (Lewis, 1986, pp. 136-192)), Lewis is able to make the claim that modality reduces to possible worlds.

So take again our possible worlds talk. Possibly $P =_{df}$ there is a possible world in which P , necessarily $P =_{df} P$ in all possible worlds. On most accounts that analysis is not reductive, because on most accounts possible worlds are in some sense modal: a possible world is a way the world might have been. For Lewis things are different. His possible worlds are concrete individuals that exist in just the way the actual world does. So it is open to him to claim that modality *just is* the distribution of truth values across the space of possible worlds.

3. The Problem

Now that we have an adequate handle on the view targeted by the problem of advanced modalizing, I want you to imagine that you yourself hold that view. That is, I want you to imagine that you yourself are a genuine modal realist. Now, given you allegiance to genuine modal realism there are going to be some things that you will accept automatically, simply as

a result of your newfound modal realism. For instance, you'll certainly think that **(Worlds)** is true.

(Worlds): there are at least two (concrete, spatiotemporally isolated...)¹
possible worlds.

The vast majority of us also take for granted the principle of possibility introduction **(PI)**². So you also will likely accept **(PI)**.

(PI) anything that is true is also possibly true.

On the basis of these two beliefs, you would be reasonably expected to accept **(◇Worlds)**, since **(◇Worlds)** follows from **(Worlds)** and **(PI)** by modus ponens.

(◇Worlds): ◇There are at least two possible worlds.

But now you're in trouble. According to you, '◇P' just means 'there is a possible world in which P'. So **(◇Worlds)** just means:

(∃Worlds): There is a possible world in which there are at least two possible worlds.

(∃Worlds) is false. There is no world that contains any other world – they are spatiotemporally isolated individuals (Lewis, 1986, pp. 69-78).

This is bad. According to the Lewisian translation scheme, **(∃Worlds)** is the counterpart theoretic translation of **(◇Worlds)**^{3,4} and they express the same proposition. But if that were true it couldn't be the case that one is false while the other is true.

¹ For brevity, the parenthetical comment will be omitted in what follows.

² Here I'm using John Divers' label (Divers, 1999, pp. 217).

³ The Lewisian translation scheme for closed sentences of the form $\diamond\phi$ works by replacing the possibility operator, '◇', with ' $\exists\beta(W\beta\&\dots)$ ', and restricting all quantifiers in ϕ to the domain of β (Lewis, 1968, p. 117). Performing that operation on **(◇Worlds)** - ' $\exists y\exists z((Wy\&Wz)\&y\neq z)$ ' - results in $(\exists x(Wx\&(\exists y\exists z((Iyx\&Izx)\&(Wy\&Wz)\&y\neq z)))$: There's something, x, which is a world, and has two things, y and z, which are also worlds, as parts, and y and z are distinct.

⁴ Lewis stuck to the 1968 translations as his preferred version, but he did offer some alternative conceptions of being in, or part of a world (Lewis, 1983, p. 40). One alternative, x gets to count as being in world w if it is *partly in* w. That is, if x shares a part with w. If we make this amendment to the translation schema, the problem remains. 'There is a possible world which shares a part with two possible worlds' must be false since, on modal realism, worlds do not overlap. On the other alternative, x gets to count as *being in* world w if x exists *from the*

John Divers (Divers, 1999) calls sentences like (\Diamond **Worlds**) ‘Advanced Modals’, hence ‘the problem of advanced modalizing’. (\Diamond **Worlds**) is not the only example of an advanced modal, there are many. Take, for example, the fusion of my left shoe with all its counterparts and call the resulting object ‘*Shoes*’. According to modal realism, it is true that *Shoes* exists. Given (PI), then, it must also be *possible* that *Shoes* exists. That is, it must be true that ‘possibly, the fusion of my left shoe with all of its counterparts exists’. But this sentence can’t be true according to modal realism since it translates into counterpart theory as ‘there is a world which has as a part, the fusion of my left shoe with all of its counterparts’ and no object is a counterpart of anything in its own world (Lewis, 1968, p. 114).⁵

4. A Solution

What to do about this? A natural response is that the classic translation scheme for sentences of modal logic is faulty. We saw, above, that the operation of the translation scheme on advanced modals has a distorting effect. This might be taken as data that shows the original translation scheme incapable of translating certain sentences of modal logic⁶.

But I think this is misguided. It’s true that the standard translation schema of genuine modal realism fails to accommodate advanced modals. But this is a *feature* of Lewis’ account, rather than a bug. Advanced modals should be taken to show that (PI) is false, and there are some sentences that are just not the kind of sentences that are apt to be modalized.

(PI) is highly intuitive and well entrenched, so I want to give some examples of analogous cases in which a (PI)-type principle is unattractive. I’ll argue that largely (PI)-type

standpoint of *w*. And *x* exists from the standpoint of *w* iff ‘it belongs to the least restricted domain that is normally – modal metaphysics being deemed abnormal – is appropriate in evaluating the truth at that world of quantifications’ (Lewis, 1983, p. 40). The sentence ‘there is a possible world such that two possible worlds exist from the standpoint of it’ is presumably false; if the least restricted normal domain of a world excludes modal metaphysics, presumably other possible worlds are excluded on this basis. (Divers, 1999, pp. 223-5).


⁵ This is a version of the example used by Hud Hudson in ‘A True, Necessary Falsehood’, 1999.

⁶ Divers (1999, 2014) opts for this strategy and gives a ‘redundancy analysis’ of advanced modals. Dorr does something similar (2010, see especially p. 4 and footnote 9). Jago (forthcoming) argues that this would be one good strategy but that it fails on independent grounds.

principles are unattractive in these cases because of the kind of analysis under question, and that Lewis' analysis of modality is this very kind of analysis itself.

Advanced Localising.

Imagine that you and I are working on the puzzle below.



We get to talking philosophically, and I suggest an analysis of what it is for something to be true everywhere in the puzzle ('everywhere-truth') vs. what it is for something to be true somewhere in the puzzle ('somewhere-truth'). My suggested analysis is as follows:

(Everywhere-Truth): P is everywhere-true iff for any piece, P is true of that piece.

(Somewhere-Truth): P is somewhere-true iff for some piece, P is true of that piece.

The sentence 'there is an eye' would, on this analysis, be an example of a somewhere-truth, since it is true of some piece of the puzzle that 'there is an eye'. It would not, on the other hand, be an example of an everywhere-truth because 'there is an eye' is not true of all

pieces. Contrastingly, ‘there is a colour’ *would* be an example of an everywhere-truth, since this sentence is true of every piece of the puzzle.

Imagine now that you present me with a problem case. ‘What, then’ you ask ‘of the sentence ‘there is a fish’? It seems to me that, although this sentence is demonstrably true, it can have neither the status of everywhere-truth (since it is not true of every piece of the puzzle) nor the status of somewhere-truth (since it is not true of any one piece of the puzzle on its own). Surely this is unacceptable since ‘there is a fish’ is demonstrably true. How can a sentence be true, but not somewhere-true? After all, claiming that a sentence is true, and not true anywhere sounds dangerously close to saying that it is true and not true...’

This problem case assumes the localizing analogue of **(PI)**, call that **(LI)** and define it as follows:

(LI): if P is true, then P is somewhere-true.

But given this analysis I offered, (LI) must be false. Since my account analyses somewhere-truth and everywhere-truth in terms of localised goings-on at particular regions, we should fully expect that ‘there is a fish’ might be true but not somewhere-true. ‘There is a fish’ is *just not the kind of thing* that can be somewhere- or everywhere-true. That may seem odd given our ordinary usage of ‘somewhere’ and ‘everywhere’, but my analysis can still be helpful in allowing us to sort puzzle pieces, and maybe even more helpful than any analysis which allows ‘there is a fish’ to be somewhere-true.

Similarly, genuine modal realism analyses modal claims in terms of how modal space as a whole plays out; what’s possible in one region depends on what’s going on in the other regions and the modal facts are reflected in the similarities and differences between possible worlds. We can perfectly legitimately make claims about what is true of the overall spread of possible worlds – claims about the *pluriverse as a whole*. But these claims are not apt to be modalized, just as advanced spatial claims are not apt to be somewhere- or everywhere-true.

Advanced Temporalizing

Imagine, now, that you and I are discussing the B-theory. I suggest to you that, given the B-theory, an important distinction among the truths is between those that are true some of the time (the sometimes-truths) and those that are true all of the time (the always-truths). ‘Fran is sitting’ and ‘something exists’ are both truths, but they are importantly different. ‘Fran is sitting’ is true only when indexed to certain times, whereas ‘something exists’ is true regardless of which time it is indexed to. Because this distinction seems intuitive and important, I decide to offer an analysis for it. That analysis goes as follows:

(Always-Truth): P is always-true iff P is true at every time slice

(Sometimes-Truth): P is sometimes-true iff P is true at some time slice.

‘There are guinea pigs’, then, is sometimes-true, since it is true of some time-slices (this one) but not true of others (those in the very distant past). Contrastingly, ‘something exists’ is always-true, since it is true of every time-slice.

But now we encounter a problem. You may think that, given the meaning of sometimes- and always-truth, any true sentence must also be sometimes-true. After all, intuitively, if x is not sometimes-true, then x is *never* true, which means that x is false. So, perhaps you believe a principle I will call ‘sometimes-truth introduction’ **(SI)**. Now, on my theory, the sentence ‘there are at least two time-slices’ is true. So, given the principle **(SI)**, ‘there are at least two times slices’ ought to be sometimes-true. But my analysis cannot allow for that. ‘There are at least two time-slices’ is not true of *any* time-slice, since no single time slice contains two time-slices.

It seems to me a perfectly acceptable response to insist that sentences of the type ‘there are at least two time-slices’ are just not the kinds of things that get to be sometimes- or always-true. This needn’t have the result that sentences like this are true but never-true; a sentence which is not apt to be sometimes-true is just as unable to be never-true.

Consider our conversation progressing as follows:

You: What about the sentence ‘there are swans’; is that sentence always-true or sometimes-true?

Me: well, since that is only true at some time-slices...

You: Oh, don’t misunderstand me. When I say ‘there are swans’ I mean to have my quantifiers range over *all* of the time slices, not just one. My sentence is something more like ‘there are swans *somewhere in space-time*’ or ‘there are (*ranging over all time-slices*) swans’

Is ‘there are (*ranging over all time-slices*) swans’ something that is true all of the time or something that is true only some of the time? To say that it is sometimes-true feels inappropriate; the overall spread of time-slices does not change, so it would seem odd to suggest that this sentence is sometimes-true and sometimes-false - in virtue of what is it false rather than true? On the other hand, it also feels inappropriate to claim that ‘there are (*ranging over all time-slices*) swans’ is always-true. ‘Something exists’ is always true because there is always, in each time-slice, something to make it true. ‘there are (*ranging over all time-slices*) swans’ is not like that. Its truth is not ‘reaffirmed’ by every time-slice. Both of these options feel wrong because it is a mistake to attribute sometimes- or always-truth to this sentence; there are some sentences that just aren’t apt to be sometimes- or always-true.

You might have very legitimate, independent reasons for wanting to reject the above analyses of advanced localizing and advanced temporalizing. But it ought not be on the grounds of their inability to analyse advanced sentences. To do so would be to reject our definition of what it is to be green because it does not accommodate the number 2. In the next section, I will spend some time bringing this final analogy out.

Advanced Modalizing

Advanced modalizing is unproblematic for genuine modal realism in the just way that advanced temporalizing and advanced spatializing are unproblematic for the above analyses.

The modal realist analyses possibility and necessity in the following way:

(Possibility): P is possible iff P is true at some possible world.

(Necessity): P is necessary iff P is true at every possible world.

As a result of this analysis, there will be some truths that are neither possible nor necessary. This is not a problem. ‘There are at least two possible worlds’ is just *not the kind of thing* that can have the predicates of possibility or necessity attached to it. This has the direct result that (PI) is false, just as (SI) and (LI) and false.

And in the case of advanced modals, we can identify the kind of sentence that is unmodalizable. A close look at the mechanics of the problem of advanced modalizing shows us that it occurs when a particular kind of sentence is scoped by a box or diamond.

Recall the discussion earlier in the paper about the role of the translation schema in the reductive component of the Modal Realist view. Lewis’ showed that we can have a language which can express any sentence expressible in quantified modal logic⁷, but without having to introduce modal operators by demonstrating that we can translate any sentence of quantified modal logic into a sentence of counterpart theory using his translation schema. The translation schema does this by replacing any possibility or necessity operator with an existential or universal quantifier (respectively) with possible worlds as its domain of quantification. So, using Lewis’ 1968 rules, $\Diamond\phi$ translates into the language of counterpart theory as $\exists x(Wx \& \phi^x)$ – *there is a thing, x, x is a world and ϕ is true at x*. Importantly, ϕ^x is itself a complex sentence. It expresses ‘ ϕ is holds in x’ and is formed by restricting the range of all quantifiers in ϕ to the domain of x. So if ϕ denotes ‘ $\exists y(Fy)$ ’ – *there is a thing, y, and y*

⁷ It can also express some sentences that quantified modal logic cannot (Lewis, 1968)

is F - then ϕ^x should be understood as expressing ' $\exists y(Iyx \& Fy)$ ' – *there is a thing, y , y is in x , and y is F* . Equivalently, if ϕ were to denote ' $\forall y(Fy)$ ' – *for all y , y is F* - then ϕ^x should be understood as expressing ' $\forall y(Iyx \supset Fy)$ ' – *for all y , if y is in x , then y is F* . In the process of translating $\diamond\phi$ into the language of counterpart theory, all quantifiers in ϕ are restricted to the domain of *some particular world*.

As a more concrete example, let's say that I want to translate the sentence 'Possibly, there are blue swans' ($\diamond(\exists yBSy)$) into the language of counterpart theory. In order to do this, the diamond is first swapped out of the formula and replaced by ' $\exists x(Wx\dots)$ '. Then the sentence that is scoped by the diamond (in this case ' $\exists yBSy$ ') is said to hold in world x , so all the quantifiers in that sentence are restricted to x . This yields the fully translated sentence ' $\exists x(Wx \& \exists y(Iyx \& BSy))$ ' (There is a world, x , and there is a blue swan, y , and y is in x). In most (perhaps all) ordinary contexts, when we make quantified claims about what is possible or necessary we are making claims about what is possible or necessary *for this world*. That is, when we say 'Possibly, there are blue swans' the quantifier 'there are' is ranging over only things in the actual world. We don't really care, when we are talking normally, about what might have been the case at some *other* possible world; we care about what might have been the case here. So, when Lewis' translation schema shifts the range of the quantifier in ' $\exists yBSy$ ' it isn't too damaging to the meaning of the sentence. After all, this is what the translation scheme is *supposed* to do: it takes claims about what *might have been* true here, and makes them into claims about what *is* true in some other world. However, there is another way that we might have understood the sentence 'Possibly, there are blue swans' that is damaged by this. If we take 'there are' in 'there are blue swans' to range more than just the actual world, then the artificial restriction is problematic. All of a sudden, the restriction of quantifiers in ϕ to the domain of some particular world makes a difference to the *meaning* of the original sentence.

Take, for example, the sentence ‘there is a guinea pig’. The quantifier in this sentence could be interpreted in many ways, two of which would be;

\mathbf{GP}^R : ‘ $\exists x(\mathbf{GP}x)$: ‘there is (implicitly restricted to the actual world) a guinea pig’

\mathbf{GP}^U : ‘ $\exists x(\mathbf{GP}x)$: ‘there is (implicitly restricted to the pluriverse) a guinea pig’

In the name of brevity, let me introduce two new pieces of terminology: call any quantifier that ranges over at least two things that are in different worlds *non-worldbound*, and call any quantifier that ranges over only things within a single world *worldbound*. The quantifier in \mathbf{GP}^R , then, is worldbound whereas the quantifier in \mathbf{GP}^U is non-worldbound. And we can see the difference between these two sentences. \mathbf{GP}^R says that there is a guinea pig in the actual world, \mathbf{GP}^U says there is a guinea pig in the pluriverse. While it is true that the latter implies the former, they do mean different things, just in the way that ‘there’s a shoe in the closet’ and ‘there’s a shoe in the house’ mean different things. \mathbf{GP}^R entails that I am spatiotemporally related to a guinea pig, \mathbf{GP}^U does not. We should expect, then, that modalized versions of \mathbf{GP}^R and \mathbf{GP}^U will also have different meanings. That is, if we stick a diamond or a box out in front of each sentence we should expect the following two different sentences to result:

$\diamond\mathbf{GP}^R$: ‘ $\diamond\exists x(\mathbf{GP}x)$ ‘Possibly, there is (implicitly restricted to the actual world) a guinea pig’

$\diamond\mathbf{GP}^U$: ‘ $\diamond\exists x(\mathbf{GP}x)$ ‘Possibly, there is (implicitly restricted to the pluriverse) a guinea pig’

However, genuine modal realism cannot appreciate the difference between these two sentences. Both $\diamond\mathbf{GP}^R$ and $\diamond\mathbf{GP}^U$ translate as $\exists\mathbf{GP}$.

$\exists\mathbf{GP}$: ‘ $\exists x(Wx \& \exists y(Iyx \& \mathbf{GP}y))$: ‘there is a possible world in which there’s a guinea pig’.

You'd be within your rights to be non-committal about what exactly $\Diamond GP^U$ means. But it's nonetheless compelling that $\Diamond GP^U$ and $\Diamond GP^R$ are different sentences. If you start off with two different sentences, the scoping each with a modal operator ought not return the same sentence in each case.

$\Diamond GP^U$ and $\Diamond GP^R$ come out of the Lewisian translation as the same sentence because the only the non-worldbound quantifier in $\Diamond GP^U$ is artificially restricted by the translation schema so that it becomes a worldbound quantifier. In other words, $\Diamond GP^U$ is incorrectly translated because it is an advanced modal. Given this understanding, we can understand 'advanced modals' as sentences in which an 'non-worldbound' quantifier is scoped by a modal operator.

The example of $\Diamond GP^U$ is instructive of two facts, i) that any modal claim will have an advanced and an ordinary reading (because any quantified claim has a worldbound and a non-worldbound reading)⁸ and ii) that some advanced modals are less obvious than others. $\Diamond GP^U$ is not obviously problematic because it translates into a truth: 'there is a world in which there is a guinea pig'. Since $\Diamond GP^U$ translates into a true sentence, we don't get the odd results we got before where some advanced modal was deemed true but it's counterpart theoretic translation was deemed false. Nonetheless, 'there is a world that contains a guinea pig' is still an unsuccessful translation of 'Possibly, there is (implicitly restricted to the pluriverse) a guinea pig', and we can see this by considering GP^U and GP^R above. The difference is that the mistranslation happens to coincide with something we think true (namely that it's possible that there are (restricted to the actual world) guinea pigs). Notice also that the real work of the problem is done by the quantifiers. The modal operators function much the same across worldbound and non-worldbound cases (indeed, that is part of

⁸ Divers' (2014) account also has this result.

the problem), and what the diamond and the box mean in each case is just the same. The difference comes from the status of the quantifiers.

To reflect this fact, I will henceforth distinguish between ‘advanced quantifications’ and ‘advanced modals’. ‘Advanced modal’ takes the definition given earlier. ‘Advanced quantifier’ should be understood as a sentence of quantified logic which includes an non-worldbound quantifier⁹.

However, you might worry that my definition rules out some perfectly good examples of advanced modalizing. Imagine that you arrive at a coffee shop and sit down. Once you have sat, you realize that your view isn’t very good and you would prefer to have a seat with an ocean view. With a view to fixing this, you ask yourself, ‘is there somewhere I could have sat such that I would have had an ocean view?’. The correct way to interpret this is as asking something like ‘is there somewhere I could have sat such that from *that spot* (the alternative possible spot) I would have had a view of the ocean?’ This is different from asking ‘is there somewhere I could have sat such that *this very spot that I occupy right now* had a view of the ocean?’. In the first sentence, you are asking whether there are other seats in the café that have views of the ocean, in the second sentence you are asking whether there is an alternative café layout which would be such that the spot you originally picked would have an ocean view. We can ask how things would have been different in a different seat, or we can ask how the café layout as a whole might have been different.

Given modal realism, we can ask similar kinds of questions about the modal space. We can ask how things would have been different in a different world, or we can ask how the layout of the pluriverse as a whole might have been different. In the latter sense, we might consider the sentence ‘possibly, there are talking donkeys *at w*’ (where *w* names a world). That is, we might consider whether the region of modal space that currently lacks talking

⁹ Parsons suggests talking of ‘advanced quantification’ rather than ‘advanced modalizing’, but I would contend that, given the interaction between the modal operators and the non-worldbound quantifiers, it is worth it to distinguish the two (Parsons, forthcoming, pp. 14-15).

donkeys (namely the region that contains world w) might have instead contained talking donkeys. This is different from asking whether there might have been talking donkeys. In asking the latter, you are essentially asking whether there is a way things could have been that includes you and blue swans being in the same world (regardless of which world that is). In asking the former, you are asking whether there is a way things could have been such that *the space of possibilities were different*.

This sentence fails the criteria in my definition, since the quantifier ‘there are’ in ‘ \Diamond there are talking donkeys at w ’ only ranges over things in w , and my definition requires that the quantifier scoped by a modal operator have a wider range than that. On the other hand, we might also think that this sentence is an advanced modal.

A closer look will reveal that my analysis *does* count this as a case of advanced modalizing. First, notice that when translating ‘ \Diamond there are talking donkeys at w ’ into counterpart theory, we would need to use two existential quantifiers to really replicate the claim: $\exists x \exists y (Wx \& TDy \& Iyx \& x=w)$ – *there are two things, x and y , x is a world and y is a talking donkey and x is in y and x is w .*¹⁰ While the second quantifier is restricted to w by ‘ Iyx ’, the first quantifier is non-worldbound. This makes ‘possibly, there are talking donkeys at w ’ an advanced modal.

You might think that we should represent ‘there are talking donkeys at w ’ by giving the sentence ‘ $\exists x(TDx)$ ’ and stipulating that the existential quantifier be restricted to the domain of w , but this would be inaccurate. ‘ $\exists x(TDx)$ ’ would yield the claim ‘there are talking donkeys *here*’, that is, it would yield the more ordinary claim. Importantly, there is a difference between ‘possibly, there are talking donkeys at w ’ and ‘possibly, there are talking donkeys’ when uttered at w . The former makes an advanced claim about how the pluriverse might have been different as a whole, whereas the latter makes an ordinary modal claim.

¹⁰ Thanks to Karen Bennett for helping me see this.

An additional virtue of this account is that it can explain why ersatz modal realisms are immune to problems of advanced modalizing. Ersatz views are actualist views, meaning that they will never countenance the use of non-worldbound quantifiers. While it is true that, according to ersatz modal realism, we can make claims about what is true at some possible world, doing so does not amount to endorsing a non-worldbound quantifier on this view, since all possible worlds are actual existents. That is, in claiming ‘there is a possible world at which there is a talking donkey’, on the ersatz view, I am making a (true) claim in which the initial quantifier ranges only over individuals that share a world.

But what exactly do I mean when I say that certain truths are *just not the kind of things* that can be modalized? My suggestion is that in cases like these a category mistake is being made. The sentence ‘the theory of relativity is sleeping’ looks false, but does not entail the truth of ‘the theory of relativity is awake’ even though we generally expect that something that fails to be asleep must be awake. This is because the theory of relativity is *just not the kind of thing* that can be asleep or awake. And that’s quite different from its being straightforwardly false that the theory of relativity is sleeping: the false sentence ‘the theory of relativity is sleeping’ has a different status than the false sentence ‘Fran is 7 feet tall’.

Something like this is going on in the case of advanced modals: advanced modals are false, but this is not to say that their negations are true. ‘Possibly, there are at least two possible worlds’ is false, but so is ‘*it is not the case that, possibly there are at least two possible worlds*’. That is to say, while it is not possible that there are at least two possible worlds, this fact does not entail that it is impossible that there are at least two possible worlds, just as its being false that the theory of relativity is sleeping does not entail that the theory of relativity is not sleeping, or that the theory of relativity is awake.

But there is an important difference between the sentence ‘the theory of relativity is sleeping’ and advanced modals such as ‘possibly, there are at least two possible worlds’. The

former is obviously and immediately ridiculous. The latter sounds like it makes complete sense. Indeed, it took me a good ten pages to explain why the latter is a strange sentence, and only a couple of sentences to explain why the former is strange. So why should we think that advanced modals are category mistakes, given that they are relatively easy to make sense of? The burden is on the modal realist to explain why this is so.

This concern about the claim that advanced modals are category mistakes might be cashed out in various different ways. One immediate thought might be that, in order to qualify as a category mistake, a sentence must be *meaningless*, and it is far from clear that advanced modals are meaningless. There are two ways that the genuine modal realist might respond to this concern. One way would be to reject the meaninglessness view of category mistakes (following Magidor, 2010). But a much more powerful way to respond would be to highlight the difference in context between the quantified sentences of ordinary conversation, and the quantified sentences of fully fledged modal realism. I will take up the latter approach.

On Ofra Magidor's understanding, category mistakes are infelicitous because they trigger presuppositions that are not acceptable in the conversational context. So, for example the predicate 'is angry' in 'the fridge is angry' triggers the presupposition that the fridge experiences emotions. But that the fridge experiences emotions cannot be accommodated in most conversational contexts – since in most cases conversants will take it for granted that the fridge does not experience emotions. The failure of the predicate 'is angry' to trigger the requisite presupposition results in an infelicitous sentence (Magidor, 2013, pp. 131-2).

Adopting this account, we should think that advanced modals also suffer from presupposition failure. In a sentence of the form 'x is necessary', the predicate 'is necessary' invokes the presupposition that quantifiers in x span a region equal to or smaller than a single possible world; that the quantifiers in x are *worldbound*. When we substitute for x some

advanced quantification, these presuppositions can no longer be accommodated and the resulting sentence suffers a presupposition failure.

Crucially, though, this presupposition failure will only occur in certain, relatively restricted, contexts. That's because ordinary contexts read quantifiers as exclusively world-bound: whatever we substitute for x , it will not be an advanced quantification. Almost all ordinary contexts are actualist contexts, they involve sentences that quantify over only objects in the actual world. Given that this is the case, the presuppositional trigger for 'x is necessary' will be accommodated with ease.

This brings out an important point. The presupposition trigger for sentences of the form 'x is necessary' will always be accommodated as long as we are actualists. The actualist simply does not have the ontological resources to make advanced quantifications, so their quantifiers will always be worldbound and their substitutions for x will always be unproblematic. It is only when we embrace the genuine modal realist's ontology that we can start to talk in advanced terms. So the presuppositional failure only kicks in when we fully commit to the ontology of the genuine modal realist. This can explain the apparent benignity of advanced modals; we are so used to assessing quantified sentences in terms of our actualist sensibilities that all quantifications are read as being world-bound, even after we make the shift to genuine modal realism.

Conclusion.

Advanced modals are sentences in which a non-worldbound quantifier is scoped by a modal operator – where a quantifier is 'non-worldbound' if any two things within its range are not spatiotemporally related to one another. This account helps to explain what is problematic about advanced modals, since the Lewisian translation schemes treats all quantifiers within the scope of the modal operator as though they were worldbound. While it

is true that advanced modals such as ‘possibly, there are at least two possible worlds’ come out false given the Lewisian translation scheme, this is not the disaster that most have taken it to be. We should regard advanced modals as committing category mistakes – they attempt to assign the properties of possibility and necessity to things that are *just not apt to be modalized*. This has the result that some truths – the advanced quantifications – are neither necessary nor possible nor even impossible. That is to say, the sentences ‘possibly, there are at least two possible worlds’, ‘necessarily, there are at least two possible worlds’ and ‘it is impossible that there are at least two possible worlds’ all return the truth-value false, even though ‘there are at least two possible worlds’ is true. This, of course, has the result that **(PI)** is false: not all true sentences are possible. But this fact should also not be worrying given the conditions under which it is false. Modal realism can keep its translation schema, and its reductive ontology, even in the face of advanced modalizing.

Bibliography:

- Bennett, K. (2005). Two Axes of Actualism. *The Philosophical Review*, 114(3), 297-362.
- Divers, J. (2014). The Modal Status of the Lewisian Analysis of Modality. *Mind*, 123(491), 861-872.
- Divers, J. (1999). A Genuine Realist theory of Advanced Modalizing. *Mind*, 108(430), 217-239.
- Dorr, C. (unpublished). How to be a Genuine Modal Realist.
- Fine, K. (2005). *Modality and Tense*. Oxford, UK: Clarendon Press.
- Hudson, H. (1999). A True, Necessary Falsehood. *Australasian Journal of Philosophy*, 77(1), 89-91.
- Jago, M. (forthcoming). Advanced Modalizing Problems. *Mind*.
- Lewis, D. K. (1986). *On the Plurality of Worlds*. Oxford, UK: Blackwell.
- Lewis, D. K. (1983). Postscripts to "Counterpart Theory and Quantified Modal Logic", In Lewis, D. K. (1983). *Philosophical Papers Volume I*, Oxford, UK: Oxford University Press.
- Lewis, D. K. (1968). Counterpart theory and Quantified Modal Logic. *The Journal of Philosophy*, 65(5), 113-126.
- Magidor, O. (2013). *Category Mistakes*. Oxford, UK: Oxford University Press.
- Marshall, D. (unpublished). A Puzzle for Modal Realism.
- Noonan, H. (2014). The Adequacy of Genuine Modal Realism. *Mind*, 0(0), 1-10.
- Parsons, J. (unpublished). Against Advanced Modalizing.
- Rosen, G. (1993). A Problem for Fictionalism about Possible Worlds. *Analysis*, 53(2), 71-81.
- Rosen, G. (1990).
- Sider, T. (2002). The Ersatz Pluriverse. *The Journal of Philosophy*, 99, 279-315.