

Defeating Extreme Modal Skepticism

Lars Enden

How do we know that something is possible? On the one hand, knowledge of possibilities seems unproblematic, because possibility claims are so weak. Even when the proverbial cat is not on the mat, it still *could have been* on the mat, right? Intuitively, we know that the cat could have been on the mat, because saying that something is *possible* is just not saying very much at all. The standards of justification for possibility claims seem to be very, very low. On the other hand, knowledge of possibility is mysterious when we realize that we do not have a very clear understanding of the truth conditions for such propositions. Isn't that odd? How can we know that it is possible that the cat is on the mat when we are not even sure what it *takes* for this claim to be true? How can we know a proposition when we are not clear about what that proposition *means*? These are the questions of the extreme modal skeptic. The extreme modal skeptic believes that we are never justified in our beliefs concerning the possibility of false propositions. Now, I have never actually met an extreme modal skeptic, but these questions are not easy to answer, and I think a close inspection of this issue will help to clear the way for some much needed progress in the epistemology and metaphysics of modality.

In this paper, I will state what I take to be a powerful argument for extreme modal skepticism, and I will work toward what I take to be a plausible response. The response will require the development of unique and original views of modal semantics, modal metaphysics, and modal epistemology.

1. The Argument for Extreme Modal Skepticism

The argument for extreme modal skepticism begins with three theses. It is important to the argument that will develop that these three theses have strong intuitive appeal; the argument is skeptical, so intuitive appeal is all it really has.

The first thesis is that the truth conditions for modal propositions are mind-independent. The possibility of the cat being on the mat is about something outside of anyone's way of thinking or conceiving of cats or mats or the relationship of on-ness. To display the intuitive appeal of this thesis it is worth considering the following. If all minds were removed from the world, is it more intuitively plausible that modal truth would also be removed or that some modal propositions would remain true? It seems more intuitively plausible that some modal propositions would remain true; remove all of the minds from the world, and it remains true that Humphrey could have won that election in 1968. So, at a first glance anyway, this thesis is more intuitively plausible than its negation. I call this thesis the semantic assumption because it assumes that the correct semantic theory of modality will provide truth conditions that are mind independent.

The second thesis is, very roughly put, that being justified in one's belief concerning any proposition with mind-independent truth conditions, modal or otherwise, requires that one have perceptual evidence that the truth conditions for the proposition have been satisfied. The idea here is that having a justified belief requires that one have a reason to believe that the truth conditions for that proposition have been satisfied. When the truth conditions in question are mind-independent, then one must be able to access that mind-independent world in some way in order to determine whether the truth conditions have been satisfied or not, and the only way we have for accessing the mind-independent world is through perception. To use a well-worn example, a blind man, due to his disability, is not able to have perceptual evidence for color claims, so it seems intuitively clear that he simply does not have the right kind of evidence to justify his beliefs concerning color.¹ Therefore, it is intuitively plausible that in order to be

¹ Except perhaps through the testimony of others. Presumably, however, such testimony can be traced back to someone's direct visual experience, and justification is transitive across this testimonial transaction.

justified in one's beliefs concerning a mind-independent reality, one must have some perceptual evidence for those beliefs. I call this thesis the epistemic assumption.

The third thesis is that one can only have perceptual evidence for the possibility of a proposition if that proposition is true. In other words, it is not possible to have perceptual evidence to justify one's belief that the truth conditions for a non-actual possibility claim have been satisfied. To bastardize a famous bit from Kant: perception can teach us what is, but it cannot teach us what might be. We cannot see possible sights or hear possible sounds, unless of course those possible sights and sounds become actualized. So, in general, it is intuitively plausible that having perceptual evidence for a possibility claim can only occur if the claim is actually true. For non-actual possibility claims, there is nothing "there" to be perceived. I call this thesis the metaphysical assumption.

It may already be fairly clear how these three theses imply extreme modal skepticism, but here is a version of the argument:

1. S is justified in believing that "Possibly P," and P is false (Assumption for reductio).
2. The truth conditions for "Possibly P" are mind-independent (Semantic Assumption).
3. If S is justified in believing that "Possibly P," and the truth conditions for "Possibly P" are mind-independent, then S has perceptual evidence that the truth conditions for "Possibly P" have been satisfied (Epistemic Assumption).
4. Hence, S has perceptual evidence that the truth conditions for "Possibly P" have been satisfied (from 1, 2 and 3).
5. If P is false, then it is impossible for S to have perceptual evidence that the truth conditions for "Possibly P" have been satisfied (Metaphysical Assumption).

6. Hence, S has no perceptual evidence that the truth conditions for “Possibly P” been satisfied (from 1 and 5).

7. Therefore, if P is false, then S is not justified in believing that “Possibly P” (from 1, 4, and 6).

The conclusion of this argument supports extreme modal skepticism: no one is ever justified in believing that a non-actual possibility is true.

The argument for extreme modal realism is similar to Paul Benacerraf’s famous dilemma in the philosophy of mathematics.² Indeed, the philosophy of mathematics and the philosophy of modality are close cousins, and, like most cousins, they rarely talk to one another even though they share many of the same family problems. However, Benacerraf’s dilemma (modified to apply to modality) is importantly different than the argument just presented. For one thing, Benacerraf was no skeptic: he wondered how we could reconcile mathematical realism with a plausible epistemology of mathematics. For another thing, Benacerraf suggested that the apparent inconsistency was between mathematical realism and a causal condition on knowledge. The argument just presented does not rely on a causal account of knowledge. The extreme modal skeptic merely suggests that some way or other of obtaining perceptual evidence be available in order for a belief concerning an external world to be justified. Still, Benacerraf’s worry is similar to the extreme modal skeptic’s worries. They both are concerned about how realism concerning apparently abstract objects can be reconciled with a commonsensical understanding of how we are justified in our beliefs concerning a world beyond our own minds.

One way that we could respond to the argument is by embracing the conclusion: we could become extreme modal skeptics by just denying that we have any knowledge of non-actual

² Paul Benacerraf, “Mathematical Truth,” in *Philosophy of Mathematics: Selected Readings*, ed. Paul Benacerraf and Hilary Putnam (Cambridge University Press, 1983), 403-420.

possibilities. But I take it that this is an extremely implausible view. Even Peter van Inwagen, who calls himself a modal skeptic, would not go so far. Van Inwagen thinks that many possibility claims *made by philosophers* are not knowable, like “zombies are possible,” but he does not deny that we know a great many possibilities that are more closely connected to the affairs of everyday life, like “the couch could have been placed two inches to the left of where it is now.”³ I suggest, then, that we do our best to avoid extreme modal skepticism and only return to it if no other solutions prove to be more plausible.

That leaves us with denying one of the three intuitively plausible theses with which we began. Each of the theses has its opponents, so none of the three are beyond criticism. I will not attempt here to catalogue the various objections that could be leveled against each of the assumptions because that would take us too far off course. Instead, I want to focus attention on the assumption that I think is least plausible and to develop a theory of modality that has the resources to provide a plausible objection to the argument for extreme modal skepticism. The thesis I think is the least plausible is the metaphysical thesis. I believe that it is possible to have perceptual evidence that the truth conditions for the possibility of a proposition have been satisfied even if that proposition is false. The theory I shall develop will show how. In order to do this, however, I will need to reconsider both the truth conditions for modal propositions and what counts as perceptual evidence.

2. The Ways of Perception and the Ways of the World

Let us first consider perceptual evidence. The paradigm example of perceptual evidence is knowledge derived from the physical senses. Although I am surely risking opening an entirely different can of worms, I propose to consider responding to the argument for extreme modal

³ Peter van Inwagen, “Modal Epistemology,” *Philosophical Studies* 92 (1998): 67-84.

skepticism by considering, in a rough way, what role perception plays in knowledge derived from the physical senses.

The story is a familiar one from introductory epistemology courses: the world appears to me (and to you) to be a certain way. I take this appearance as evidence that the world is more or less the way it appears. Whether, and how, I am justified in jumping this infamous gap is of course a matter of considerable debate. Nevertheless, I want to point out that it is easy and natural to use “way”-talk when describing the basic intuitive mechanics of perception and the gathering of perceptual evidence: one perceives the *way* the world appears to be, and one uses this perceptual data as evidence that the world is more or less the *way* that it appears. In fact, it is difficult even to talk about perception and knowledge of the external world without using way-talk.

Interestingly, way-talk also plays a starring role in metaphysical discussions concerning modality. Consider David Lewis’s famous argument for possible worlds.

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’.⁴ (Emphasis mine)

Lewis ultimately suggests that possible worlds are worlds in the same sense that the actual world is a world. The actual world is “I and all of my surroundings,” and the other possible worlds are things of the same sort that are spacio-temporally isolated from one another. But as Robert Stalnaker has pointed out, “If possible worlds are ways that things might have been, then the

⁴ David Lewis, “Possible Worlds,” in *The Possible and the Actual: Readings in the Metaphysics of Modality*, (Cornell University Press, 1979), 182.

actual world ought to be *the way things are* rather than *I and all my surroundings*. *The way things are* is a property or state of the world, not the world itself.”⁵ So, like the classical philosophical discussions about the role perception plays in knowledge of the external world, the debate over the ontological status of possible worlds also turns (partially) on how to make sense of our intuitive way-talk.

What would happen, then, if we take way-talk seriously? What if we suppose that *ways*—whatever those are—are the objects of perception as well as the objects of modal truth? The theory that I want to develop begins with this basic idea. Fortunately, others have come before me, so some work has already been done in this direction.

3. Nature Realism

As already noted, Stalnaker suggests that *the way the world is* and the *ways that the world could have been* are properties of the world, but he doesn’t think that these *ways* can be further analyzed. Peter Forrest, however, disagrees. Forrest has developed Stalnaker’s rough idea into a more developed version of what has become known as nature realism.⁶ Forrest suggests that for every Lewisian possible world, there is what he calls a world-nature, by which he means “the property which is the conjunction of all its (natural) non-relational properties.”⁷ Therefore, the actual world-nature, according to Forrest’s theory, is the conjunction of all of the natural, non-relational properties that the actual world instantiates. But Forrest says that a single set of natural, non-relational properties and relations can be conjoined in many different *ways* each of which forms a different property, some of which may be world-natures and others of which may not be. So, for Forrest, there is more to the actual world-nature than just the conjunction of all of

⁵ Robert C. Stalnaker, “Possible Worlds,” in *The Possible and the Actual: Readings in the Metaphysics of Modality*, (Cornell University Press, 1979), 228.

⁶ The name was given by John Divers. See his *Possible Worlds* (London: Routledge, 2002), 177.

⁷ Peter Forrest, “Ways Worlds Could Be,” *Australasian Journal of Philosophy* 64 (1986): 15.

the natural, non-relational properties of the actual world; these properties must also be conjoined in the right way. Therefore, Forrest's complete theory of modality will also require a theory of structural composition for properties together with a means for determining which of these structural compositions counts as a world-nature and which do not.

According to Forrest, world-natures are composed of natural, simple properties and relations. The properties and relations are simple in the sense that they are not reducible to other properties and relations. Although Forrest does not explain in detail what he means by a natural property, presumably he has in mind something like Lewis's notion of natural properties and relations as the properties and relations that "carve at the joints" of reality.⁸ Forrest writes, "*Vermillion*, or rather some determinate shade of vermilion, might turn out to be a natural property. So might *being a proton*. But *being grue* and similar hybrids are not."⁹ Presumably, *being grue* is not a natural property because it fails to "carve at the joints" of reality, and, presumably, *being (some determinate shade of) vermilion* is a natural property because it does "carve at the joints."

The meat of Forrest's theory is what he calls primitive operations of composition. These are intended to explain the way in which the simple, natural properties and relations can be composed into more complex properties and relations in a structural (non-mereological) way.

The first operation Forrest calls *taking the product*. Forrest explains it as follows.

If R is an m -adic property or relation and S is an n -adic one, then $R \times S$ is the $(m+n)$ -adic relation which holds between $x_1, \dots, x_m, y_1, \dots, y_n$ just in case R holds between x_1, \dots, x_m and S holds between y_1, \dots, y_n .¹⁰

⁸ See David Lewis, "New Work for a Theory of Universals," *Australasian Journal of Philosophy*, 61 (1983): 343-77.

⁹ Peter Forrest, "Ways Worlds Could Be," 15, fn 5.

¹⁰ *Ibid.*, 17.

To use Forrest's example, suppose that G is some determinate shade of green and that H is some determinate shade of blue. The product $G \times H$ will be the determinate relation of color contrast that holds between x and y just in case x is that determinate shade of green and y is that determinate shade of blue.¹¹ The product of any two properties results in a relation—that unique relation that holds between two objects whenever any they instantiate the two properties. It should be carefully noted, however, that the product operation is not symmetric; $R \times S$ is not the same relation as $S \times R$, but its converse.

The next operation is contraction. The idea behind contraction is that two or more relata in a single relation can be identified with one another.¹² Again to use one of Forrest's examples, *knowing oneself* is a property, but it is the contraction of the *knowing* relation. It is a contraction in the sense that in the relation *knowing*, the two places of the relata are identified with one another to produce the property *self-knowing*. In other words, *S knows herself* can be understood as *S knows S* with the two relata *contracted* into one. More generally, to form a contraction of a relation, one just identifies two or more of the relata with one another.

The last operation is projection. Forrest writes, "Consider an n -adic relation R . Suppose a_1, \dots, a_n are related by R . Then as a consequence, the sum of $a_1 + \dots + a_n$ has a property, namely being the sum of parts related by R . I call this the (monadic) *projection* of R ."¹³ The (monadic) projection of a complex *relation* is the *property* of being the mereological sum of parts related by that relation. For example, suppose that B is the property of *being a ball*, and R is the property of *being red*. The product $B \times R$ is the relation that holds between any ball and any red thing. By contraction, the two relata are identified, and the projection of the resulting contraction is the property of *being a red ball*.

¹¹ *Ibid.*

¹² *Ibid.*, 18.

¹³ *Ibid.* Italics in original.

To see how all of these operations work together to demonstrate structural composition, I will add a few more symbols to Forrest's. First, to express that a relation is n -adic, I will use n different numbers in subscript after the capital letter representing the relation. For example, O_1 might be the property *being an oxygen atom*, and $B_{1,2}$ might be the dyadic relation *being bonded to*. The point of the numbers is not just to show how many relata are in the relation, but they also prove to be a convenient way to express the operation of contraction, which should be apparent from the following examples.

Step 1: Start with *natural* properties and relations (ones that “carve at the joints”).

O_1 is the natural property *being an oxygen atom*.

C_1 is the natural property *being a carbon atom*.

$B_{1,2}$ is the natural dyadic relation *being bonded to*.

Step 2: *Form a product* of the natural properties and relations.

$O_1 \times O_2 \times C_3 \times B_{4,5} \times B_{6,7}$

(The seven different numbers show that this is a heptadic relation.)

Step 3: *Contract* some (or all) of the relata.

$O_1 \times O_2 \times C_3 \times B_{1,3} \times B_{2,3}$

(Changing the subscript numbers indicates that relata with the same number are identical. In this example, the relation has been contracted from a heptadic relation to a triadic one.)

Step 4: *Project* the resulting product into a single property.

The sum $(1+2+3)$ has the property *being carbon dioxide* (O—C—O).

To see how different properties can be composed of these same parts, consider an alternative to Steps 3 and 4.

Step 3a: *Contract* some (or all) of the relata.

$$O_1 \times O_2 \times C_3 \times B_{1,2} \times B_{2,3}$$

(The only difference between this and the previous Step 3 is that the contraction has been done in a slightly different way—the second and fifth relata in the heptadic relation have been identified instead of the third and fifth relata.)

Step 4a: *Project* the resulting product into a single property.

The mereological sum (1+2+3) has the property *being (unstable) C—O—O*.

By contracting the products of natural properties in various ways, we can demonstrate how the same properties and relations can be structurally composed in different ways. The very same properties and relations are involved in both *being carbon dioxide* and in *being (unstable) C—O—O*, but the *way* they are composed is not. Forrest's operations give us a means by which to demonstrate such structural differences in complex properties and relations.

Forrest needs more than this, though, because not all structural properties that can be defined by the operations are world-natures. For example, the projection of the property $C_1 \times S_1$, where C_1 is the property *being cubic* and S_1 is the property *being spheroid*, had better not be a world-nature, since it is clearly not possible for any single thing to have the property *being both cubic and spheroid*. Forrest needs a way to divide the structural properties into world-natures and non-world-natures.

According to Forrest, world-natures are those properties that are capable of instantiation by themselves. This, by itself, is of course not an adequate way to make the needed distinction, since it includes the modally loaded word “capable.” So Forrest needs a modally neutral way of indicating which structural properties are capable of instantiation by themselves and which are not. Forrest's solution is to introduce the concept of a *completion*. The completion of a property,

P, is the property *having no other properties other than P and itself*. As an example, *being blue* is not a world-nature (cannot exist by itself) because it does not have a completion: *being blue and nothing else* is not a natural property since anything that is blue must, at least, also have some shape or other. Still, *being blue* is a conjunct of many complex properties that have completions, like, say, the property *being a blue box that weighs 1 pound, has square sides, has a hinged lid, has...etc*. Impossible properties, like *being completely blue and completely red*, have no completions and are not conjuncts of any complex properties that have completions. So, for Forrest, a world-nature is simply a completion—any completion. Properties that are not completions are not world-natures.

Forrest needs one more bit of machinery to make his theory work. Lewis’s semantics for modal claims requires a relation between propositions and worlds called “true at” (or sometimes “true in”). For example, P is possible, according to Lewis, if and only if, there is at least one possible world such that P is *true at* that world. For Forrest’s theory to offer a plausible alternative to Lewis’s theory, Forrest must replace this relationship because there are no worlds for Forrest’s propositions to be *true at*. Forrest’s solution is simple. He merely changes the terminology from *true at* to *true under*.

- (i) ‘Possibly P’ is true iff P is true under some world-nature; and
- (ii) ‘Necessarily P’ is true iff P is true under all world-natures.

But of course, if *true under* is going to do the work that Forrest requires, it had better get the actual world-property—the way *this* world is—right. Therefore,

- (iii) ‘Actually P’ is true iff P is true.¹⁴

Nature realism is starting to take shape. It began with Stalnaker’s idea that modal realism can be maintained without regarding possible worlds as the same kind of thing as the actual

¹⁴ *Ibid.*, 16.

world. We can think of them as *ways the world could be* rather than as full-blooded worlds à la Lewis. Forrest further developed Stalnaker's theory by showing how we can understand world-natures as complex structural properties, suggesting the concept of a completion, and providing a modal semantics.

4. A Problem for Nature Realism

There is, however, a serious problem for Forrest's version of nature realism. Forrest claims that an advantage of his theory over Lewis's is that it makes a categorical distinction between the actual and the possible, the actual is particular and the merely possible is universal.¹⁵ This is in keeping with the actualist spirit of nature realism as presented by Stalnaker. However, Forrest builds this idea into his semantics. According to his semantics, "Possibly P" is true just in case P is true under some world-nature, and "Actually P" is true just in case P is true.¹⁶ The problem is that if this semantic theory is correct, then what are we to make of certain obvious modal-logical principles, in particular, "Actually P implies Possibly P?"

To see the problem, it is worth noting that the actual world-nature makes no appearance in Forrest's semantic theory. One would expect Forrest's theory to say "Actually P" is true iff P is true under the actual world-nature, but instead it says "Actually P" is true iff P is true. Since Forrest believes that actuality is particular and possibility is universal, this seems to imply that "Actually P" is true in a different way than "Possibly P." In other words, being true is a different property from being true *under* the actual world-nature. Here is where the problem lies. If Forrest's semantic theory is correct, then "Possibly P" does not obviously follow from "Actually P," since it does not follow from "Actually P" that P is true under at least one world-nature. Nothing in the semantic theory suggests that a true proposition must be true *under* some world

¹⁵ *Ibid.*, 22.

¹⁶ *Ibid.*, 16.

nature. Presumably, Forrest thinks that any proposition that is true is true under the nature of the actual world, but his semantic theory makes no mention at all of the nature of the actual world.

So Forrest will have to accept this thesis:

P is true iff P is true under the actual world-nature.

Although the simplest way for Forrest to accommodate this needed thesis into his metaphysical theory would be to maintain that being true just *is* being true under the actual world-nature, he does not have to say this. Still, he must maintain that being true exactly coincides with being true under the actual world-nature, which may just be a brute fact about truth and the relation *truth under*.

Now Forrest gets the result that “Actually P implies Possibly P.” This addition to his semantic theory may seem innocuous. I am sure that Forrest himself thought it was too obvious to bother mentioning, but what I want to show now is that this clarification shows that Forrest’s theory is implausible.

To showcase the problem for Forrest’s theory, we must think a little harder about the actual world-nature. The actual world-nature is understood to be a structural property, and all the other world-natures are structural properties as well. What distinguishes the actual world-nature from the non-actual world-natures is that the actual world-nature is instantiated and the others are not instantiated. All of these world-natures exist, whether they are instantiated or not, and, since this is an actualist view, they are all part of the actual world.

Now here comes the trouble. It turns out that there are no false contingencies: any proposition that is false will be necessarily so. If this is right, then extreme modal skepticism follows from this theory because it is not possible to know something that is not true. If there are no false contingencies, then it is not possible to know one. Here is the proof that Forrest’s theory

implies that there are no false contingencies. Consider, for *reductio*, a proposition that is possible but false. Since P is false, it is not true under the actual world-nature according to Forrest's theory, but, because it is possible, it is true under some non-actual world-nature. It has already been shown that Forrest has to accept the thesis, "P is true iff P is true under the actual world-nature." The result now is that "Possibly P" has to be true under the actual world-nature, which means that the proposition "P is true under a non-actual world-nature" is itself true under the actual world-nature. But this leads to a serious problem: P turns out to be true under the actual world-nature after all, which contradicts our supposition that P is false.

Let me fill in the details of the argument. The non-actual world-natures exist as part of the actual world, but the actual world-nature is supposed to be the structural property instantiated by the actual world. Since all of the world-natures exist as part of the actual world, part of what the actual world-nature instantiates is all of these world-natures (including itself, somehow). It seems that the only way this makes any clear sense is if those other world-natures are themselves included within the actual world-nature. The other world-natures have to be somehow *represented* by the actual-world nature because the non-actual world-natures exist as part of the actual world, and the actual world-nature is instantiated by the actual world. Presumably, in order for "P is true under a non-actual world-nature" to be true under the actual word-nature, there must be some sort of correspondence between "P is true under a non-actual world-nature" and the actual world-nature. For such a correspondence to make any sense, the actual world-nature must be able to represent all of the other world-natures. This is because it is the one and only world-nature instantiated by the actual world, which includes all of the other world-natures as a part. So, anything that is true under any world-nature is true under the actual world-nature.

But wait. One might wonder whether there is a kind of representation of other world-natures within the actual world-nature short of full inclusion that might resolve this issue. But, it is not easy to see how this could be. The needed stand-ins for world-properties would have to be structural properties, because that is all that the actual world-nature is composed of, and they would have to represent somehow non-actual world-natures, which are also structural properties. So what we are looking for are structural properties that can represent structural properties. I think it is fairly clear that this only makes sense if each non-actual world-nature is simply included within the actual world-nature. It is hard to see how anything else will do. The needed entities would have to be structural properties complete enough to represent other structural properties, but not so complete as to be identical with them. Those seem like unlikely heroes. Therefore, if Forrest's theory is correct, then all of the non-actual world-natures exist as part of the actual world-nature.

To return to the main line of argument, we assumed, for *reductio* that P is false but possible. This leads to the contradictory result that P is both true and false under the actual world-nature. It is false under the actual world nature simply because it is false, and it is true under the actual world nature because to be possible is just to be true under a world-nature that is part of the actual world-nature, and, therefore, true under the actual world-nature after all. Consequently, Forrest's theory has the implausible result that there are no false contingencies; any proposition that is possible (true under at least one world-nature) must also be actual (true under the actual world-nature).

One way a nature realist could go to resolve this issue is to give up on actualism—become a possibilist. In other words, a nature realist could agree with Lewis that the *way the*

world is is identical to the world itself.¹⁷ If there is no distinction between the actual world and the actual world-nature, then the non-actual world-natures are free, so to speak, to live alongside the actual world-nature rather than being a part of it. The difference here from Lewis's ultimate view is that the possibilist nature realist emphasizes the way over the world. They both agree that the way and the world are identical, but Lewis says that they are both "I and all of my surroundings," and the possibilist nature realist says that they are both a structural property. This becomes a possibilist view rather than an actualist view because it turns out under this view that there are things that exist separate from the actual world, the non-actual world-natures.

However, if one goes this route, it is hard to see the advantage this view has over Lewis's brand of possibilism. The whole point of actualism was to try to maintain the intuitively compelling idea that the actual world is the only world. Once that idea is abandoned, there seems to be no reason not to go all in on Lewisian possible worlds. The bigger problem, however, is that we lose any chance of responding to the argument for extreme modal skepticism. One cannot have perceptual evidence that anything is true under a non-actual world-nature if these world-natures exist completely separately from the actual world. We would have no way to "check" them to see what is true under them. So it seems that denying actualism is not a plausible solution to the problem because what has led us to nature realism in the first place is a desire to respond to the argument for extreme modal skepticism, and this desire will be thwarted if we turn nature realism into a possibilist view.

5. Structural Modal Realism

I want to suggest a different solution. One that takes nature realism in a completely different direction entirely; a direction I think changes the view so drastically that it becomes a different, but related, theory altogether. I call the resulting view structural modal realism. The

¹⁷ David Lewis, *On the Plurality of Worlds* (Malden, MA: Blackwell, 1986), 86.

idea is simple; deny the existence of non-actual world-natures. There is just the world and the way that the world is. That's it. This is the view that I think can offer a plausible response to the argument for extreme modal skepticism.

Under structural modal realism, there are no possible worlds, no ways that the world could have been, no non-actual world-natures. Instead there is just the world and the way that it is, the actual world-nature. Instead of all of the non-actual world-natures, modal truth is accounted for by a proper part of the actual world-nature, what I call the necessary world-nature. The necessary world-nature is that part of the actual world-nature that could not have been otherwise. According to this view, there is no such thing as a *way that the world might have been*. Instead there is only the *way that the world is*, and part of this is the *way that the world must be*. Possibilities are simply those propositions that are not ruled out by the *way that the world must be*.

Now this may seem like a drastic departure from contemporary modal metaphysics, which may be troubling because we know the exceptional explanatory power of possible worlds. Well, first, I think it offers the only plausible realist response to extreme modal skepticism, which I will return to shortly, but I also believe that the departure is not as drastic as it might appear. To demonstrate this, consider again Lewis's argument for the existence of possible worlds.

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'.¹⁸

¹⁸ David Lewis, "Possible Worlds," 182.

Lewis argued in favor of possible worlds by touting their power to explain ordinary beliefs concerning modality. I claim at least as much for structural modal realism. I can argue in much the same way as Lewis:

It is uncontroversially true that some things must be the way that they are. I believe, and so do you, that there are countless ways that things could not have gone differently. But what does this mean? Ordinary language permits the paraphrase: there is a way that things must be in addition to the way that they are. On the face of it, this sentence is an existential quantification. It says that there exists two entities of a certain description, to wit “the way things are” and “the way things must be.” I believe that things could not have been different in countless ways; I believe permissible paraphrases of what I believe; taking the paraphrase at its face value, I therefore believe in the existence of two entities that might be called “the way things are” and “the way things must be.” I prefer to call them “the actual way” and “the necessary way.”

I, therefore, claim the same ordinary-language and intuitive benefits that Lewis claims. So structural modal realism is not as much of a departure from a commonsensical understanding of modality as it may appear.

I must now fill in the details of structural modal realism. I begin with semantics. According to structural modal realism, it turns out that there is no such thing as a non-actual world-nature. Instead what accounts for modality is the actual world-nature, which is naturally divided in two parts, the necessary and the non-necessary. Possibilities are simply understood to be those propositions that are not ruled out by the necessary part. Therefore, instead of possible world semantics, we have necessary world-nature semantics:

“Necessarily P” is true iff P is true under the necessary world-nature.

“Possibly P” is true iff not-P is not true under the necessary world-nature.

“Actually P” is true iff P is true under the actual world nature.

The necessary world-nature is understood to be part of the actual world-nature. It is that part of the actual world-nature that could not have been otherwise.

Still it is hard to see how we can have a reason to believe that a given proposition is true under the necessary part of the actual world-nature. To show how this may be done, I need to introduce some mental analogues to the actual world-nature and the necessary world-nature. I will call the mental analogues way-views. I doubly hesitate to call them “world-views” for fear either that they will be confused for certain social constructs involving biases or that they will be construed as possible-world-like. They are neither. Way-views are way-like; they are structurally composed of mental analogues to properties, which I will call, with some trepidation, sense impressions. A way-view, then, is a structure of sense impressions analogous to world-natures.

The rough idea behind a way-view is that it is essentially the way the world seems to be relative to a person. Each of us has a way-view that is constantly in flux as we gain further sensory impressions. To further extend the analogy, a proposition can be accepted or not accepted under a way-view. If S believes that P, then S accepts P under S’s way-view. In other words, part of what it is for S to believe a proposition is for it to seem to S that P is true under the actual world-nature. S’s way-view is what S takes the actual world-nature to be like; it is S’s mental representation of the actual world-nature. This is a rough sketch of course, but I hope the idea is clear enough. The most important thing to note about way-views is that they are structures in much the same way that world-natures are structures, and they are constructed in a similar fashion to the way that Forrest explained the construction of world-natures.

6. Denying the Metaphysical Assumption

Now we are in a position to see how structural modal realism provides a reason to deny the metaphysical assumption. Let us return to the basic idea of perceptual evidence that led us down this path. The story is that the world appears to me (and to you) to be a certain way, and I take this appearance as evidence that the world is more or less the way that it appears. Now that

structural modal realism has been developed, we can now understand this to mean: I have a way-view, and I take this way-view as evidence that the actual world-nature is more or less accurately represented by my way-view. In other words, our way-views *are* our perceptual evidence. It is not infallible evidence by any means, of course, but it is evidence all the same. Again, it is not clear how we jump the gap from the way the world appears to the way that the world is, but intuitively we manage this feat somehow. What I want to show is that *if* we manage this feat for knowledge of the external world, *then* structural modal realism shows that there is not really an additional problem of knowledge of non-actual possibilities. If our way-views are perceptual evidence for actuality, then they are also perceptual evidence for possibility.

Here is how it works. When S considers a particular non-actual possibility, like “the cat is on the mat” (even though it is not), S takes his or her way-view and rearranges parts of it—S mentally “puts” the cat on the mat. S reconstructs the way-view so that the cat-impression finds a place “on” the mat-impression. When S finds that he or she can construct such a way-view, S has a reason to believe that the analogous reconstruction is not ruled out by the necessary world-nature. S, therefore, is justified in believing that “the cat is on the mat” is indeed a possibility. I submit that this process of reconstructing one’s way-view counts as perceptual evidence that the proposition in question is possible. Compare this with an attempt to reconstruct a way-view to make it such that an object is both cubic and spheroid at the same time. The needed reconstruction does not result in a way-view—one is simply unable to perform the needed reconstruction. Now this might not really count as perceptual evidence for anything because failing to make the needed reconstruction may just be due to a lack of imagination. So I would not say that failing to reconstruct one’s way-view is evidence that the world-nature could not be reconstructed in the analogous way. However, when we manage to do the needed reconstruction,

then I think that counts as very strong evidence that the reconstruction in question is not ruled out by the necessary world-nature.

This is a very rough sketch of the complete theory of course, but I hope that the main idea is fairly clear. When one manages to reconstruct one's way-view to form another way-view, this provides some perceptual evidence (since one's way-view *is* one's perceptual evidence) that the actual world-nature could be rearranged in that same way. In other words, reconstructing one's way-view to form a representation of the truth conditions for P is perceptual evidence that the negation of P is not true under the necessary world-nature. If I am right about this, then it is possible to have perceptual evidence that a proposition is possible even if that proposition is false, and, therefore, the metaphysical assumption is false, and extreme modal realism is defeated.