Plenitude and Recombination

Alastair Wilson
University of Birmingham & Monash University
a.j.wilson@bham.ac.uk

10th June 2021

This is a draft of a chapter that has been accepted for publication by Oxford University Press in "David Lewis: A Historical Perspective", eds. Helen Beebee and Anthony Fisher.

Abstract: In *On the Plurality of Worlds* (Lewis 1986), David Lewis imposes a condition on realist theories of modality which he calls 'plenitude'. Lewis apparently assigns this condition considerable importance, and uses it to motivate his Humean principle of recombination, but he never says exactly what plenitude amounts to. This chapter first sets aside some obvious ways of reconstructing the plenitude criterion which do not fit with the textual evidence. An objection to modal realism due to John Divers and Joseph Melia (Divers and Melia 2002) is diagnosed as equivocating between an overly-demanding plenitude constraint and a weaker constraint which fails to establish their conclusion. An alternative deflationary interpretation of the plenitude condition has it following from an application of standard theoretical virtues to a modal realist's total theory; Lewis' correspondence provides new evidence in support of this interpretation. The deflationary plenitude criterion also has broader application, beyond Lewisian modal realism.

1. Introduction

David Lewis devotes an entire section of *On the Plurality of Worlds* (Lewis 1986; henceforth, *OPW*) to a condition which he calls *plenitude*. Lewis tells us that we need 'a way to say . . . that there are possibilities enough, and no gaps in logical space' (1986: 89). Uncharacteristically, he does not give a clear statement of the plenitude criterion, despite appearing to treat it as an important methodological constraint on theories like his own modal realism. Instead he gives us a little intellectual autobiography, relating how he at first thought of plenitude as the requirement that the following conditions be satisfied:

(1) Absolutely every way that a world could possibly be is a way that some world is.

(2) Absolutely every way that a part of a world could possibly be is a way that some part of some world is. (1986: 86)

The problem with (1) and (2) (first pointed out by Peter van Inwagen (1980)) is that, once we make the modal realist's proposed theoretical identification between worlds and ways things could be, (1) and (2) are rendered trivial. As Lewis puts it:

> (1) says only that every world is identical to some world. That would be true if there were only seventeen worlds, or one, or none. It says nothing at all about abundance and completeness. (1986: 86)

> (1), and likewise (2), cannot be salvaged as principles of plenitude. Let them go trivial. Then we need a new way to say what (1) and (2) seemed to say: that there are possibilities enough, and no gaps in logical space. (1986: 87)

Lewis's positive claims about plenitude in *OPW* are—perhaps intentionally—vague. All he says is that plenitude requires 'abundance and completeness' of worlds, and that we need a guarantee 'that there are possibilities enough, and no gaps in logical space'.

What is all this meant to mean? Does modal realism satisfy plenitude? And why care about the criterion in the first place? Answering these three questions—interpretive, metaphysical, and methodological—is the project of this paper. After a survey of the textual evidence, including Lewis's recently-published correspondence, I look in §2 at some natural ways of reconstructing plenitude, and argue that none of them fit with what Lewis says about the criterion. John Divers and Joseph Melia have argued that modal realism fails to satisfy plenitude; in §3–5 I diagnose their argument as equivocating between an overly-demanding constraint and a more acceptable constraint which fails to establish their conclusion. In §6 I propose my own interpretation of plenitude, according to which it follows from the application of standard criteria on theory-choice to the modal realist context. Consequently, analogues of the plenitude criterion can be formulated in the context of any type of theory of modality (§7). §8 is a conclusion.

2. Lewis on Plenitude

A starting point in untangling Lewis's discussion of plenitude in *OPW* is to streamline the terminology. His presentation strongly suggests that the requirement that there be 'possibilities enough' is the same requirement as 'abundance', and that the requirement that there be 'no gaps in logical space' is the same requirement as 'completeness'. In the two extracts just quoted, Lewis moves seamlessly from the one pair of requirements to the other,

giving no indication that the subject is being changed; and it would seem somewhat bizarre if plenitude amounted to a conjunction of four distinct conditions. In what follows I will assume that these identifications were intended by Lewis.

Nevertheless, the requirements we are left with are still problematically unclear. 'Abundant' is a vague term in ordinary language, and cashing out 'abundance' as the requirement that there be 'possibilities enough' invites the question: enough for what? Likewise, 'completeness' is underspecified, and 'no gaps in logical space' is metaphorical. Such imprecision is unacceptable, given that Lewis apparently intends plenitude to be a substantive feature of his modal realism. We need a clearer statement of the requirements imposed by plenitude in order to assess whether it is a well-motivated constraint on candidate analyses of modality.

How might we reconstruct Lewis's conditions more precisely? Start with 'abundance', the requirement that there be 'possibilities enough'.[1] Lewis certainly thinks that no worlds, one world, and seventeen worlds would not be possibilities enough. So perhaps abundance is the requirement that there be, say, at least denumerably many worlds. By itself, this is insufficient to ensure the variety of possibilities that he envisages. A denumerable set of worlds which contained only qualitative duplicates of the actual world would be unable to provide plausible truth-conditions for our modal talk of the sort that Lewis is after. What we need is that the following requirement be satisfied:

> ***Diversity:*** A modal realist theory should entail that the members of the set of worlds manifest sufficient qualitative diversity.

Of course, this condition is still underspecified: how much diversity, and in what respects, counts as sufficient? The requirement of Diversity is just a placeholder for a more refined

[1] It's fairly clear that 'possibilities' here means 'maximal qualitative possibilities', rather than the notion of 'possibilities' which Lewis discusses in his 1986: §4.4, where a possibility is a world plus an assignment of *de re* representation. One world may include infinitely many possibilities in the latter sense, and perhaps these infinitely many possibilities might be 'enough' for some purpose: but Lewis intends the notion of abundance to rule out there being exactly one world. His discussion of plenitude also comes long before the discussion of haecceitism which motivates his distinction between worlds and possibilities.

statement of the plenitude criterion. But rather than resolving the vagueness in the criterion any further, Lewis moves straight to introduce a principle designed to satisfy it.

The *principle of recombination*, Lewis says, can 'express' (87) or 'capture' (92) plenitude:

> Roughly speaking, the principle is that anything can coexist with anything else, at least provided they occupy different spatio-temporal positions. Likewise, anything can fail to coexist with anything else. Thus if there could be a dragon, and there could be a unicorn, but there couldn't be a dragon and a unicorn side by side, that would be an unacceptable violation of plenitude. And if there could be a talking head contiguous to the rest of a human body, but there couldn't be a talking head separate from the rest of a human body, that too would be a failure of plenitude. (1986: 88)

Motivated by concerns about conceptual circularity,[2] Lewis transformed the above explicitly modal statement of the principle of recombination into the following principle about worlds (which is free of modal vocabulary): any two parts of any two worlds have duplicates which occur together in some world or another, and any two parts of one world have duplicates which occur on their own in some distinct pair of worlds. Lewis informally[3] stated this non-modal form of the principle as follows:

> ***PR:*** 'patching together parts of different possible worlds yields another possible world.' (1986: 87–8) [4]

PR harmonizes with Lewis' Humean outlook; when introducing it, Lewis made explicit appeal to Hume's 'denial of necessary connections between distinct existences'.[5]

---

[2] These concerns may not be well-founded. Cameron (2012) argues that there is no need to formulate the principle of recombination in non-modal terms, as long as the modality involved is itself amenable to the Lewisian analysis. I return to this point briefly in §7.

[3] The vagueness of 'patching together' can be precisified by invoking patterns of perfectly natural external relations, so that any parts of any world stand in any pattern of perfectly natural external relations (e.g. Lewis 1987, 2009; see also Bricker 1996). A further refinement of PR might enable it to handle determinable properties; see §6 for more discussion.

[4] Lewis inserts a proviso, 'size and shape permitting' (1986: 89) for technical reasons relating to the limitations of a single space-time. This proviso will not be directly relevant to the arguments at hand.

[5] Lewis 1986: 87; see also Lewis 1983a: 314.

PR provides us with an important clue in interpreting plenitude. Whatever the criterion of plenitude amounts to, it must be a criterion which the truth of PR goes at least some way towards helping to meet. The simplest way in which the truth of PR could help to meet the plenitude criterion would be for PR itself to be constitutive of plenitude. Accordingly, one possible interpretation of plenitude simply identifies plenitude with the requirement that the postulates of modal realism should entail that PR is true:

Combinatorial Plenitude: A modal realist theory should entail PR.

Since recombination applied to the actual world immediately generates denumerably many worlds, with great variety amongst them, the suggestion that plenitude just is Combinatorial Plenitude has some initial plausibility. But this would be a bad interpretation of Lewis, for two immediate reasons.

Firstly, it makes a mystery of why Lewis introduced the term 'plenitude' and glossed it in terms of 'abundance' and 'completeness'. It would have been far simpler to title section 1.8 of *OPW* 'Recombination', and to insist on PR directly as a criterion of adequacy on modal realist theories of modality.

Secondly, Lewis is explicit in *OPW* that PR, applied to the actual world, fails to ensure that plenitude is satisfied:

> We can't get the alien possibilities just by re-arranging non-alien ones. Thus our principle of recombination falls short of capturing all the plenitude of possibilities. (1986: 92)

This is clear evidence against the interpretation of plenitude as Combinatorial Plenitude. PR might be necessary for the plenitude condition to be satisfied, but it is not sufficient; *a fortiori* satisfaction of PR cannot be constitutive of the plenitude criterion.

Taking a step back, we can distinguish two theoretical roles that recombination principles might be playing in Lewisian modal realism. The first role for recombination—call it the *independence role*—is to express a Humean denial of necessary connections between distinct existences. To play the independence role, the principle of recombination simply has to be true of whichever worlds there in fact are. The second role for recombination—call it the *generative role*—is to characterize a theoretically interesting set of worlds when applied to some basic elements. To play the generative role, recombination needs a suitably rich set of basic elements to work with. The passage just quoted suggests that Lewis conceived of plenitude as involving the generative role for recombination. This impression is reinforced when we

turn from *OPW* to consider some remarks on plenitude from Lewis's correspondence, in which Lewis explicitly considers candidates for supplementary principles of plenitude.

Clear evidence that Lewis though recombination insufficient to ensure plenitude comes from Lewis's letter to van Inwagen of 16 Feb 1983 (Lewis 1983a). In this letter Lewis first sketches the puzzle of plenitude, framing it in terms very similar to those he subsequently used in §1.8 of *OPW*. He then identifies a couple of ways in which recombination falls short of specifying the full plenitude of possibilities. First, he doesn't want to apply recombination to obtain any conclusions about dimensionalities of possible space-times:

> I admit [a world with 5-dimensional space-time], but wouldn't want to be driven to it just by Recombination. (1983a: 313)

This suggests the need for principles other than recombination to ensure a suitable plenitude of dimensitonalities of possible spacetimes. Second, and perhaps more decisively, Lewis argues that recombination will never capture the full plenitude of *alien* possibilities:

> . . . Recombination doesn't give us the properties that contain really alien kinds of things . . . . So Recombination doesn't wholly capture the abundance and variety of the worlds. (1983a: 314)[6]

A third limitation of recombination is identified in a letter from Lewis to Hazen of 19 April 1983; it pertains to the need for a plenitude of possibilities for determinate values of determinable properties:

> Maybe [some world that is a model of theory T] has a value, of some magnitude, that happens never to show up in [some other world that is a model of theory T]—as it were, a missing shade of blue. (1983c: 319–20)

When *OPW* is read alongside Lewis's correspondence, it becomes evident that he regarded plenitude as a substantive criterion which could not be satisfied simply by including

---

[6] This letter continues: 'Does it capture the appropriate sense of completeness, that Peter especially needs and I also want? There, I think maybe so' (1983a: 314). Here, Lewis is talking specifically about a rationalist argument due to Peter Unger which aims to establish that reality has no arbitrary features. Presumably Lewis's idea is that PR by itself ensures that there is no arbitrariness in the possibilities concerning the distribution of any particular quality, even though it fails to ensure that there is no arbitrariness in the possibilities concerning which qualities there are.

PR as part of the modal realist theory. Instead Lewis consistently looked beyond PR and canvassed other principles of a broadly generative character in order to help satisfy the plenitude criterion. We will return in §6 to the question of the methodological function of these principles. First, it is time to turn to a prominent recent objection to the Lewisian modal realist project: that no combination of generative principles available to modal realists can successfully satisfy plenitude.

3. Divers and Melia on Completeness

As part of their axiomatic reconstruction of modal realism, John Divers and Joseph Melia (2002) explicated plenitude as the requirement[7] that the theory should ensure that the following condition is met:

> ***Completeness***: [The set of all worlds contains] worlds of sufficiently many different types[8] to represent all of the possibilities. (2002: 18)

Completeness is presented as one of two corollaries of a methodological criterion which Divers and Melia call 'Accuracy'. The other corollary is as follows:

> ***Consistency***: [The set of all worlds contains] no worlds of any type that represents any impossibility. (2002: 18)

What is required of a theory if it is to jointly satisfy completeness and consistency? On a straightforward reading of the principles, they require respectively that for any true proposition of the form $\Diamond p$ there must be a world according to which it is the case that p, and that for every true proposition of the form $\neg\Diamond q$ there must be no world according to which it is the case that q. These requirements, however, go not at all beyond the trivial (1) and (2) which Lewis rejects as candidates to characterize plenitude.

---

[7] In fact, Divers and Melia take completeness to be a replacement for plenitude which improves on it by not carrying an implication that only cardinality of worlds is at stake. However, Lewis clearly did not intend any such implication, and the terminological shift seems to me to be unmotivated.

[8] In their formulation of these conditions, Divers and Melia make reference to types of world, rather than to worlds, because of a desire to allow for the possibility of duplicate worlds, about which Lewis was agnostic. For simplicity, I will ignore duplicate worlds; but nothing much hangs on this choice.

Completeness and consistency appear to jointly amount only to the statement Divers and Melia call 'P':

> ***P:*** ◊p iff there is a world according to which it is the case that p. (2002: 17)

P is part of the modal realist analysis of modality,[9] rather than a substantive condition on the adequacy of that analysis or of a wider theory incorporating that analysis. The truth of P is compatible with there being exactly seventeen worlds, for example.

Did Divers and Melia miss this point? No; in the context in which completeness and consistency are presented, 'representing a possibility' seems to be intended to amount to representing something possible 'by pre-theoretic lights' (2002: 19). So Divers and Melia in fact interpret plenitude as the requirement that the modal realist ontology should include worlds which, via the biconditional P, secure the truth of all (or most) modal propositions that we pre-theoretically judge to be true and the falsity of all (or most) modal propositions which we pre-theoretically judge to be false. This conception of the burden on the modal realist is captured in a methodological criterion which Divers and Melia dub 'Accuracy' and take to be 'one of the more general criteria of adequacy in analysis':

> (i) that the ontological component of [Lewisian Modal Realism] generates a set of worlds that determines the truth-values of the existential claims about worlds that figure as the right sides of instances of P and
>
> (ii) that these truth-values should match the truth-values on the left-side possibility claims, these truth-values assigned – by and large – on the basis of our prior modal beliefs. (2002: 18)

On this reading, Completeness and Consistency jointly express the requirement of Accuracy: a modal realist theory should entail truth-values for all modal propositions and that these truth-values should ('by and large') match the truth-values pre-theoretically assigned to them.

---

[9] In the light of issues surrounding 'advanced modalizing' (see Divers 1999, Parsons MS, Dorr MS), we must interpret 'according to which' as 'such that, at that world' rather than as an operator which restricts quantifiers in its scope to the world in question. Divers and Melia note this subtlety, but pay it no attention; it will be of no consequence for the arguments of this paper.

There is some initial plausibility to the idea that Accuracy can provide the substantive aspect of plenitude which is lacking from Lewis's (1) and (2), and from Divers and Melia's P. Our pre-theoretic modal beliefs do, it seems, allow for a very large number of (maximal) possibilities, and for a very high degree of diversity among them. If we read 'sufficiently' in the requirement of Diversity as meaning 'sufficient to (by and large) validate our prior modal beliefs', then Accuracy looks like it might provide a plausible candidate for a non-trivial version of plenitude.

As stated, however, Accuracy seems to be too strong a requirement. As well as limiting deviation from pre-theoretic modal opinion, the criterion of Accuracy requires that an acceptable modal realism should *generate* a set of worlds corresponding to this very large number of very diverse possibilities. Although it is unclear exactly what Divers and Melia intend by 'generates' it seems plausible that a necessary condition for some principles to generate a set of worlds is for the truth of those principles to entail the existence of the worlds in the set. Merely entailing that some worlds or other exist is not sufficient to generate a set of worlds in the intended sense.

We can extract from this discussion a coherent interpretation of plenitude which is certainly a non-trivial requirement:

> ***Exhaustive Plenitude:*** The postulates of modal realism ought to entail the truth-values of all modal propositions, and these truth-values ought to agree (by and large) with pre-theoretic modal opinion.

Exhaustive Plenitude is a very demanding criterion; it requires both that an adequate modal realism must fix the contents of modal reality completely, and that the modal reality so fixed must be such as to vindicate most of our pre-theoretic modal opinion. Exactly how demanding it is depends on the notion of entailment that we take it to involve.

Divers and Melia explicitly distance their notion of completeness from metalogical completeness, asserting that 'we certainly do not expect [modal realism] to be able to prove every single modal truth' (2002: 19). But, notice that this stops short of being a denial of Exaustive Plenitude; the notion of entailment involved in Exhaustive Plenitude need not involve provability. And I have not been able to find any literal way of reading Accuracy

according to which it does not impose the requirement of Exhaustive Plenitude; nor have Alexander Paseau or Ross Cameron.[10]

What Divers and Melia seem to have in mind is a reading of 'generates' such that the postulates of modal realism can generate a set of worlds with a particular property without entailing exactly which worlds are included in the set. Even if we set aside the infelicity of this usage of 'generates', there is a further problem. The difficulty is that the property that they require the set to have—that it should vindicate ('by and large') our pre-theoretic modal opinions—itself entails a great deal about which worlds are included in the set. This difficulty may not be insurmountable: perhaps there is some gap between (on the one hand) entailing that there are enough worlds to ground most pre-theoretic possibilities and few enough worlds which ground any pre-theoretic impossibilities, and (on the other hand) entailing every single modal truth. In particular, there would be such a gap if the principles of modal realism did not entail either the truth or the falsity of any modal proposition other than those about which we have no particular pre-theoretic opinion.

If we focus on agreement with pre-theoretic modal opinion and ignore the problematic implications of their term 'generates', then we can extract from Divers and Melia's discussion a condition that is much weaker than Exhaustive Plenitude. Rather than requiring that modal realism entail all the modal truths, this condition requires that it entail all of (or most of, given the 'by and large' qualification) the modal truths which are part of pre-theoretic modal opinion:

> ***Conservative Plenitude:*** The postulates of modal realism ought to *a priori* entail the truth-values of (most of) those modal propositions on the truth of which pre-theoretic modal opinion delivers a decisive verdict, and these truth-values ought to agree (by and large) with pre-theoretic modal opinion.

This formulation seems to capture the way in Divers and Melia apply Accuracy. For example:

> We follow Lewis in thinking it extremely plausible that there could have been alien natural properties. But once we allow this possibility, we can see that GMR, in its current formulation, is incomplete. . . . . [These considerations] suggest that

[10] See Paseau 2006 and Cameron 2012.

> the genuine modal realist needs to supplement the ontological component of his theory with an axiom that entails the existence of alien properties. (2002: 27–8)

In this passage, Divers and Melia point out that pre-theoretic modal opinion allows for perfectly natural properties not instantiated at the actual world, and they use this to argue that modal realism must entail that there are worlds containing such properties. This looks like a straightforward application of Conservative Plenitude.

Conservative Plenitude is a *prima facie* plausible interpretation of Lewis's plenitude criterion. It performs several of the functions of plenitude, in ensuring that possibilities are many and various. It rules out all worlds being duplicates of one another, since we pre-theoretically believe that things could have turned out differently. It rules out there being exactly seventeen worlds, since we pre-theoretically believe that there are more than seventeen distinct ways things could have turned out. And it can potentially be motivated through general criteria of adequacy in analysis. So it seems worthy of further consideration.

Our analysis of Divers and Melia's requirement of Accuracy has delivered two candidate interpretations of plenitude: Exhaustive Plenitude and Conservative Plenitude. It is not totally clear from Divers and Melia's discussion which (if either) of the resulting criteria they endorse. However, the choice between the two criteria has crucial consequences for their overall argument. Divers and Melia's goal is to argue that modal realism fails to be reductive: they claim that the only way to ensure that it satisfies the criterion of completeness is to add postulates which are explicitly modal and which thereby undermine the reductive project. To this end, they first argue that (their axiomatization of) modal realism does not satisfy Completeness. This argument can be reconstructed as follows:

1. If the postulates of modal realism are true of a set of worlds, then that set is complete. (Supposed for *reductio*.)
2. If the postulates of modal realism are true of a set of worlds S, then they are also true of a set of worlds S* which lacks all the worlds in S containing some arbitrary property P* that is not instantiated at the actual world. (Premise.)
3. S represents a possibility which S* does not: the possibility that something is P*. (From 2.)
4. The postulates of modal realism are true of S*, but S* is not complete. (From 3.)
5. *Reductio*. (From 1 & 4.)

This argument has been challenged in various ways; see Paseau 2006 and Cameron 2012 for prominent objections. I will explain later in this section why I do not find it convincing. But for now, it will help us pin down what Divers and Melia intend by their criterion of Accuracy.

Suppose that Accuracy were intended to impose the requirement of Exhaustive Plenitude. Then the argument just given would be superfluous. A much simpler argument would be available for the conclusion that modal realism fails to meet the criterion of Exhaustive Plenitude, starting from the assumption (which Divers and Melia grant, and which I will not question here) that we cannot refer to particular alien properties. That assumption yields the following shorter argument:

6. For any property P not instantiated at the actual world, the postulates of modal realism do not either entail that it is possible that something is P or entail that it is not possible that something is P. (Premise.)
7. If modal realism satisfies Exhaustive Plenitude, then for any property P the postulates of modal realism either entail that it is possible that something is P or entail that it is not possible that something is P. (Definition of Exhaustive Plenitude.)
8. Modal realism does not satisfy Exhaustive Plenitude. (From 6 & 7.)

It is obvious that modal realism does not satisfy Exhaustive Plenitude. Given that Divers and Melia take it to be non-obvious whether or not modal realism satisfies Accuracy, they cannot intend for Accuracy to be interpreted as imposing the criterion of Exhaustive Plenitude.

Since the interpretation of their term 'completeness' as imposing the requirement of Exhaustive Plenitude fails to make sense of the dialectic of Divers and Melia's paper, I will not consider this interpretation any further. I therefore turn to Conservative Plenitude. I will argue that, if we take Accuracy as imposing the requirement of Conservative Plenitude, then Divers and Melia's central argument becomes invalid.

Recall that Accuracy is supposed to factor into two components: completeness and consistency. Whether we interpret Accuracy as Exhaustive Plenitude or as Conservative Plenitude, we can derive the following principle:

> ***C-Completeness:*** If pre-theoretic modal opinion entails that ◊p, then (by and large) there is a world according to which p.

The difference between Exhaustive Plenitude and Conservative Plenitude is not captured by C-Completeness; it lies instead in the requirement, imposed by Exhaustive Plenitude but not by Conservative Plenitude, that modal realism should entail truth-values for *all* modal propositions.

How does Divers and Melia's argument fare if we replace their completeness criterion with C-Completeness?

9. If pre-theoretic modal opinion entails that ◊p, then[11] the postulates of modal realism entail that there is a world according to which p. (Supposed for *reductio*.)

10. If the postulates of modal realism are true of a set of worlds S, then they are also true of a set of worlds S* which lacks all the worlds in S containing some arbitrary property P* that is not instantiated at the actual world. (Premise.)

11. S represents a possibility which S* does not: the possibility that something is P*. (From 10.)

12. The postulates of modal realism do not *a priori* entail that there is a world according to which something is P*. (From 11.)

13. *Reductio.* (From 9 & 12.)

The revised argument is invalid. (9) and (12) are perfectly consistent with one another, as long as pre-theoretic modal opinion does not deliver a decisive verdict on whether it is possible that something is P*. And it does not deliver any such verdict.

Pre-theoretic modal opinion (we grant for the sake of argument) entails that there is no limit on the number of possible perfectly natural properties non instantiated at the actual world. But pre-theoretic modal opinion does not entail anything about specific perfectly natural properties not instantiated at the actual world. How could it? We could not be in a

[11] I am dropping the 'by and large' qualification to simplify the argument. But we can run the argument any number of times, making use each time of some new cardinality of properties in place of the single property P* used in the main argument. This, we may suppose, leads to enough distinct violations of pre-theoretic modal opinion to be incompatible with agreement 'by and large'.

position to refer to any such property in our modal talk, or to grasp any such property in our modal thought.[12] It follows that in order to capture C-Completeness, the postulates of modal realism need not distinguish between S and S*. To be valid, what the argument needs is an additional premise such as the following:

14. For any proposition p, either the postulates of modal realism entail that ◊p or the postulates of modal realism entail that ¬◊p.

(14) is entailed by Exhaustive Plenitude, but not by Conservative Plenitude.

I conclude that Divers and Melia's objection to modal realism fails. It seems to me that whatever plausibility it does have derives from an equivocation on 'completeness'. For completeness to be a plausible constraint on modal realism, we have to understand it as motivated by Conservative Plenitude. But so construed, the inference from (3) to (4) fails. This inference goes through if we understand completeness as motivated by Exhaustive Plenitude, since Exhaustive Plenitude, via (14), licenses the move from 'S is complete' to 'any set which lacks members of S is incomplete'; but then the premises of Divers and Melia's argument are implausibly strong, and the argument is in any case redundant. On either interpretation of Accuracy, then, their argument is unpersuasive.

I have argued that modal realism cannot satisfy Exhaustive Plenitude; but I do not think the modal realist should be worried by this conclusion. In the next section, I argue that Exhaustive Plenitude is an implausibly strong condition on modal realist theories, and that it was not what Lewis intended by his talk of plenitude.

4. Against Exhaustive Plenitude

Exhaustive Plenitude was not a constraint that Lewis himself accepted. He responds (1996) to an argument that his theory's underspecification of the nature of modal reality leads to contradiction, defending the 'modal mystery' resulting from his modal realism and rejecting

[12] I set aside any cases of descriptive reference to alien properties: for example, via a description like 'the alien property which most closely resembles charge'. Such reference requires descriptive resources richer than anything Lewis acknowledged. If preferred, the argument can be reformulated in terms of alien properties with minimal resemblance to any actually instantiated properties.

a principle offered by Maudlin (1996) which leads to contradiction when combined with such modal mystery. It is unclear that Lewis even saw the phenomenon of modal mystery as any kind of cost of modal realism. In the following passage from *OPW*, he implies that it would be hubristic to think that we could settle all modal questions, even in the idealized limit of inquiry:

> On still other questions, there seems to be no way at all of fixing our modal opinions, and we just have to confess our irremediable ignorance. I think one question of this kind concerns natural properties. Is it absolutely impossible for one particle to be both positively and negatively charged? Or are the two properties exclusive only under the contingent laws of nature that actually obtain? I do not see how we can make up our minds; or what guarantee we have that there must be some way to settle the question. Certainly we are not entitled just to make the truth be one way or the other by declaration. Whatever the truth may be, it isn't up to us. (1986; 114)

Lewis sees no threat here to the reductive ambition of modal realism, implying that he did not accept Exhaustive Plenitude as a constraint on theories of modality.

Exhaustive Plenitude is doubtful for reasons extending beyond the specific case of modal realism. Theories can be coherent and informative about some subject matter without entailing the truth-value of every proposition about that subject matter. A useful comparison, drawn by Cameron (2012), is with proposed analyses of 'right action' in meta-ethics. Act-consequentialism says that the right action is the one which maximises total wellbeing. But act-consequentialism is silent on the question of which specific actions maximise total wellbeing. By itself, this is no objection whatever to act-consequentialism; we raise problems for the view only by arguing that in some specific situation S, act A maximises total wellbeing, but that nonetheless A is not the right action to perform in S. Act-consequentialism tells us what rightness consists in for actions: it does not tell us which particular actions are right. Likewise, we should only require of modal realism that it tell us what necessity and possibility consist in for propositions. Modal realism need not go on to specify which particular propositions are possible and which impossible.

We can make similar comparisons from elsewhere in what Huw Price and Frank Jackson (1997) call the 'M-worlds'. An account of the metaphysics of mathematics needs to say what it is for some mathematical entity to exist, or for some mathematical statement to be true.

But it does not need to say exactly which mathematical entities exist, or exactly which mathematical statements are true. That is a matter for mathematicians, not for metaphysicians. Likewise, an account of the metaphysics of mind needs to say what it is for some creature to be in some mental state, and perhaps what it is for some mental state to be conscious, perceptual, etc. But it does not need to say exactly which creatures are in which mental states, or exactly which mental states are conscious, perceptual, etc. That is a matter for psychologists, not for metaphysicians. I take it that considerations of this sort, applied to the modal case, undermine the requirement of Exhaustive Plenitude.

We can usefully distinguish the project of analysing modality in modal realist terms from the project of providing a satisfactory total theory which incorporates modal realism. Adopting a realist analysis of modality need not involve taking on significant commitments about the extension of modal reality. All the analysis tells us is that where goes modal reality, there go the modal facts. Of course, we want more from an analysis of modality than that it should have the right logical form. We want some reason to adopt a total theory incorporating it. So even after rejecting Exhaustive Plenitude as a constraint on modal realist theories of modality, we may still expect modal realists to tell us various informative things about what modal reality is like, and hence about what is possible and necessary. Ross Cameron puts the point clearly:

> An account of what possibility is is one thing, an account of what is in fact possible another, and these two things shouldn't be confused. Lewis owes a story about the extent of worlds, certainly; and the resulting account of what is possible had better not be too revisionary with respect to our pre-theoretic modal beliefs; but there should be no demand that this story about the extent of worlds fall out from the analysis . . . (2012: 10)

The debt Cameron thinks Lewis owes us—a story about the extent of worlds which is not too revisionary with respect to our pre-theoretic modal beliefs—is in the close vicinity of Conservative Plenitude. In any case, the requirement of Conservative Plenitude looks like a reasonable one to impose. Is it a good interpretation of Lewis' plenitude criterion?

5. Lewis on Conservative Plenitude

In the light of the foregoing, Conservative Plenitude looks like a promising interpretation of plenitude. Alas, Lewis explicitly rejects this way of giving content to the plenitude requirement:

> We might read (1) as saying that every way we *think* a world could possibly be is a way that some world is; that is, every seemingly possible description or conception of a world does fit some world. Now we have made (1) into a genuine principle of plenitude. But an unacceptable one. So understood, 1) indiscriminately endorses offhand opinion about what is possible. (1986: 87)
>
> . . . for a thesis of my position I'd like something that doesn't just endorse unspecified modal opinion. Besides, I'm not so gullible as to think that all of us are wholly right about what ways a world could possibly be! (1983a: 313)

Lewis here draws an implicit contrast between 'offhand opinion about what is possible' and informed, reflective opinion about what is possible. If plenitude goes with either of these, it goes with informed, reflective opinion. I take this to show that preservation of our pre-theoretic modal beliefs, or of our 'modal intuitions', does not exhaust what Lewis took to be the requirements of plenitude. There must be some other component of plenitude, satisfaction of which requires substantive investigation or reflection, and which is sometimes capable of over-riding our pre-theoretic modal beliefs.

Plenitude seems to be a matter of a theory's consequences for the nature of the plurality. But it does not require that every detail about the plurality should be entailed. Nor does it require that the details about the plurality entailed should always match pre-theoretic modal opinion, or even ideally reflective modal opinion. Where does this leave us?

6. Plenitude principles as laws of the plurality

It fits well with Lewis' moderately naturalistic methodology to apply the same criteria of theoretical goodness to theories in metaphysics that we apply to theories within theoretical science. (Whatever you may think of this sort of epistemological naturalism in metaphysics, it is undoubtedly a prominent theme throughout *On the Plurality of Worlds*.) Here is a clear statement:

> As best we can, I think by seeking a theory that will be systematic and devoid of arbitrariness, we arrive at a conception of what there is altogether: the possible worlds, the possible individuals that are their parts, and the mathematical objects . . . . (Lewis 1986: 111)

In the light of this general manifesto, one potential interpretation of plenitude is that it is not a *sui generis* requirement on modal realist theories, but is rather an application of some general criterion (or criteria) of theory choice to the specific question of which worlds there are.

Lewis's general account of theory-choice criteria in science is offered as part of his 'best-system analysis' conception of the laws of nature as (roughly) the set of true universal generalizations *about an individual world* which strikes the best balance between simplicity and strength.[13] But in the context of modal realism, the motivation for restricting the best-system account to apply only to individual worlds is unclear, and in apparent tension with methodological naturalism. There is no obvious reason why the best-system account must be restricted only to theories of the goings-on within individual worlds, and a natural thought is that the account might be applied to the modal realist theory itself. Doing so yields a set of true universal generalizations *about the whole Lewisian plurality* which strikes the best balance between simplicity and strength;[14] call these *laws of the plurality*.

According to the general Lewisian methodology I have described, a good theory incorporates a strong and simple set of laws. Accordingly, a good theory of modal reality will incorporate a strong and simple set of laws of the plurality. Could these laws, or some subset of these laws, correspond to 'a genuine principle of plenitude' of the kind which Lewis is seeking? Indeed, I think we can extract at least two candidate laws of the plurality from his discussion, both of which it seems reasonable to call principles of plenitude. The prospect of more than one principle of plenitude is in line with Lewis' usage: to my knowledge he never talks of 'the principle of plenitude', or otherwise implies uniqueness. And in another

---

[13] See Lewis 1973: 72–7 for the original statement of the best-system analysis of lawhood. I set aside the extension to chancy laws which Lewis later proposed; laws of the plurality are presumably non-chancy.

[14] In the quote just given, Lewis talks of maximizing systematicity and non-arbitrariness instead of maximizing strength and simplicity. It seems likely that this is just a terminological variation on the theoretical virtues appealed to in the best-system analysis, rather than a substantive difference of doctrine.

letter to van Inwagen, Lewis talks of upholding 'principles according to which the pluriverse doesn't have unduly arbitrary-seeming features' and adds '[r]ecombination is such a principle, but it doesn't do the whole job' (1983d: 331).

Lewis's principle of recombination plainly scores well on both of the Lewisian criteria for lawhood. It is highly informative about the nature of modal reality. And it assures us that there will be no arbitrariness resulting from some world *w* containing some entities in some spatial arrangement, but no world containing duplicates of those entities in a spatial arrangement which differs only in acceptable ways from that at *w*. Hence the principle of recombination achieves an impressive balance between simplicity and strength. It surely counts as a principle of plenitude, if anything does; but as we have seen, Lewis explicitly states that recombination 'falls short of capturing all the plenitude of possibilities' (1986: 92).

Although in *OPW* Lewis offers no additional candidate principle of plenitude, context strongly suggests that the remaining component of plenitude involves an abundance of possible perfectly natural properties, including alien properties. This impression is confirmed by Lewis's correspondence from the period leading up to the writing of *OPW*, in which he states—albeit in negative form—a principle about modal reality which acknowledges its alien parts:

> You ask how I express the 'seeming content' to the effect that the plenitude of worlds outruns what you can get by recombination of the elements of actuality. How about:
>
> > Some parts of some worlds are wholly alien to our world; that is, they are not duplicates of any parts of our world, nor are they divisible in such a way that all the parts into which they are divided are duplicates of parts of our world. (1983b: 318)

One element which seems to be missing from this proposed formulation is a specification of cardinality—just seventeen perfectly natural alien properties would fulfil the above condition, for example. A cardinality condition would be straightforward to add to provide a more informative principle in the same vein, although the resulting principle might then be much harder—perhaps impossible—for us to know.

Hence I would propose that the Lewisian laws of the plurality are (at least) twofold: that recombination holds true of the entire plurality, and that there are infinitely many distinct

perfectly natural properties, *including infinitely many which are alien to our actual world*. The former is the principle of recombination; the latter we might call the *principle of miscellaneity*. I will not here attempt to make the principle of miscellaneity precise: that is a task for those who would elaborate specific modal realist worldviews. Divers and Melia have a relevant suggestion—their 'alien postulate' (2002: 30). But beyond that, the problem of formulating a suitable principle of miscellaneity seems to have been unjustly neglected by modal realists.

One might think that the principle of miscellaneity as stated above looks a poor candidate for a law of the plurality: it makes reference to a particular individual, the actual world, but we do not expect laws (of any variety) to refer to particular individuals (especially not to arbitrarily chosen particular individuals). This feature can potentially be removed by requiring that the principle of miscellaneity hold at all worlds, such that every world has infinitely many perfectly natural properties alien to it. It's not clear, though, what Lewis would have made of this proposal, since his given reason in *OPW* for endorsing alien properties is that there is 'no reason to think we are privileged to inhabit' an 'especially rich world' (1986; 92)—he does not argue that there is no especially rich world in the first place.

Whether we formulate a principle of miscellaneity by specific reference to the actual world or not, problems may remain in ensuring rich enough property structures. Specifying the cardinality of the set of perfectly natural properties, even alien ones, still might not ensure abundance of the relevant kind: depending on how we exactly we choose to regiment determinate vs determinable properties within a broadly Lewisian system, it might be possible to meet the letter of any cardinality requirement simply with enough different determinates of one single determinable.[15] Unsurprisingly, Lewis was alive to this concern, and makes an initial proposal aimed at resolving it in a letter to Hazen:

> OK; Recombination needs to be supplemented by some sort of principle of Inter/Extrapolation. (1983c: 319–20)

There is some room for debate about whether such a principle of inter/extrapolation would be best seen as a supplement to PR, or whether the possibility of intermediary determinates might instead be derived from a suitably sophisticated recombination principle. That

[15] Thanks to Alex Roberts for discussion here.

question is likely to turn on delicate issues in the metaphysics of quantities; we can leave it open here.

Together with facts about the existence and nature of the actual world, the principle of recombination, the principle of miscellaneity, and a principle of inter/extrapolation seem sufficient to entail all of our uncontroversial pre-theoretic modal opinions. At least, they seem sufficient to entail all of the pre-theoretic modal opinions which motivate Divers and Melia. Perhaps, then, joint satisfaction of these three principles is constitutive of Lewis's requirement of plenitude.

On the interpretation I am proposing, the plenitude condition itself amounts to the requirement that modal realism should include a principle of recombination, a principle of miscellaneity and a principle of inter/extrapolation. This three-pronged requirement may then be understood as following from a more general methodological principle in the vicinity which I will call Plenitude*: modal realism should incorporate a plausible set of laws of the plurality, where lawhood is determined by simplicity and strength according to the same criteria as best-system laws of individual worlds.

> ***Plenitude****: A modal realist theory should entail a collectively plausible candidate set of laws of the plurality.

This interpretation of Lewis as motivated to satisfy Plenitude* fits well with much that Lewis says about plenitude. To begin with, it makes good sense of the following passage:

> A world to which no individuals, worlds or properties are alien would be an especially rich world. There is no reason to think we are privileged to inhabit such a world. Therefore any acceptable account of possibility must make provision for alien possibilities... We can't get the alien possibilities just by rearranging non-alien ones. Thus our principle of recombination falls short of capturing all the plenitude of possibilities. (Lewis 1986: 92)

Lewis here tells us that 'capturing' plenitude requires us to 'make provision for' alien properties. The reason given is *not* that alien properties are intuitively possible—it is that we have no reason to think that our world includes all of the properties that there are. This suggests that plenitude includes the requirement that the ontology delivered by a theory of modality should be *non-arbitrary*. And this is captured by the simplicity condition built into the notion of laws of the plurality. Candidate laws which made specific reference to a specific

world or set of worlds—those containing all and only the same perfectly natural properties as the actual world—would not be as simple as candidate laws lacking any such reference.

Further evidence that Lewis thought of plenitude in terms of laws of the plurality comes from the following passage:

> Although recombination will not generate alien worlds out of the parts of this world, it nevertheless applies to alien worlds. It rules out that there should be only a few alien worlds. If there are some, there are many more. (Lewis 1986: 92)

What is striking about these passages is that the restrictions imposed by plenitude in each case are conditional in form. They say that *if* there are worlds of a certain kind, *then* there are other worlds which can be generated[16] from those worlds by certain acceptable transformations. If world A and world B only differ according to an acceptable transformation, then a theory of modality which allowed that A but not B exists would be arbitrary in the sense to be avoided.

What an acceptable transformation is will depend on how we cash out the details of a modal realist theory. For Lewis, changing the spatial arrangements of any objects within a world is an acceptable transformation; so is duplicating any object in any world any number of times; so is any combination of these procedures. So too, it turns out, is combining duplicates of objects from distinct worlds in any spatial arrangement. In each case, the motivation for thinking that these transformations is acceptable is that there is nothing about any particular spatial arrangement A of entities E which could provide a reason why A is possible while a simple rearrangement of duplicates of E is not. (Whether we think that this motivation is a good one, of course, depends on what we think of Lewis's Humean agenda.)

The account of principles of plenitude as laws of the plurality makes sense of the conditionals invoked in these passages. In virtue of their strength, universal generalizations are natural candidates for lawhood. But universal generalizations entail specific conditionals: the universal generalization 'all Fs are G's' entails, for every object O, that if O is F then O

[16] I am interpreting Lewis' use of 'generated' here as metaphorical, rather than as signalling any sympathy for an anti-Humean interpretation of PR. The Humean account of lawhood seems to be an even better match for laws of the plurality than it is for more familiar laws of nature.

is G. In the light of this, any candidate principles of plenitude will entail a whole host of indicative conditionals: they will say that if worlds of some kind exist, then other worlds of that kind also exist. Failures of such conditionals, where they involve acceptable transformations, would lead to the 'gaps in logical space' which Lewis described plenitude as ruling out.

It is interesting to consider the epistemological status of laws of the plurality. Lewis's response to the question of how we acquire knowledge of the plurality of worlds is as ingenious as it is frustrating. He suggests that we can infer that particular other concrete worlds exist by combining the modal realist analysis with our pre-theoretic knowledge of what is and is not possible. This epistemological line, if cogent at all, can yield knowledge of the laws of the plurality. In order to allow for some fallibility in our pre-theoretic modal opinions, we may need to appeal to an inference to the best explanation in the move from the claim that recombination is roughly true to the claim that it is universally true. But if inferences to the best explanation of this sort are admissible in ordinary scientific contexts, as they appear to be, then epistemological naturalism suggests that they are equally admissible when it comes to the laws of the plurality.

The Lewisian line of response to the epistemological argument has been widely criticized.[17] But it is difficult to pin down the exact problem with it. One thought is that, for us to know there are enough worlds via this route, we would have to know that modal realism is correct; but we are not in a position to know modal realism is correct absent some reason to believe that recombination holds true. This line of response is developed by Cameron (2007); call it the epistemological-circle objection. If the epistemological-circle objection is successful, then the route to plenitude from the combination of the modal realist analysis with our prior modal beliefs is unavailable. But this is no objection to my claim that Plenitude* is a good interpretation of Lewis on plenitude. Since Lewis evidently thought his modal epistemology adequate to deliver knowledge about the worlds taken individually, he would have thought it equally adequate to deliver knowledge about the worlds taken collectively; that is, he would have thought his modal epistemology adequate to ground knowledge of laws of the plurality.

---

[17] See e.g. Chihara 1998 and references therein.

The conception of plenitude that I have suggested is in a sense deflationary. Plenitude* is not a *sui generis* requirement on modal realist theories of modality; rather, it amounts to the application of general criteria on theory-choice which are motivated by reasons extending beyond the modal realist context. As such, on this interpretation plenitude needs no special methodological pleading—which makes sense of why Lewis does not engage in any.

7. Plenitude beyond Lewisian Modal Realism

Plenitude* extends straightforwardly to non-Lewisian versions of modal realism, for example to the Everettian modal realism defended in Wilson 2020 or to the inflationary-cosmology modal realism described in Knobe et al. 2006. Each of these theories comes with its own candidate for laws of the plurality. In Everettian modal realism, the Schrödinger equation and the initial quantum state play the role of the law of the plurality; in inflationary-cosmology modal realism, the laws of the plurality are the inflationary dynamics and the initial state.

But Plenitude* can also be extended to apply to non-modal-realist theories. Whatever (non-eliminativist) sort of theory of modality we prefer—ersatzist, neo-conventionalist, or modalist—we can formulate an analogue of the Plenitude* criterion. Adequate versions of each kind of theory ought to incorporate strong and simple general principles concerning what is and is not possible (while agreeing, by and large, with pre-theoretic modal opinion). This gives us:

> ***Generalized Plenitude*:*** A theory of modality should incorporate plausible general principles concerning what is and is not possible, and these principles should provide an optimal balance between simplicity and strength.

Generalized Plenitude* is a constraint on a theory of modality that is orthogonal to the general character of the worlds. So Generalized Plenitude* is neutral on whether worlds are concrete individuals, or sets of sentences, or fictional individuals, or any other sort of thing.

We might try to use Generalized Plenitude* as the basis for an argument in favour of modal realism. One of the main advantages that modal realists (whether Lewisian, Everettian, or inflationary-cosmological) claim over their ersatzist opponents is the reductive nature of their theories: modal realism promises an analysis of modality in non-modal terms. Other theories require us to take some modality as primitive, for example via a primitive relation of

consistency between propositions. If Generalized Plenitude* is a legitimate requirement on theories of modality, then theories which take modality as primitive will need to include modal notions in the formulation of principles concerning what is and is not possible.

The modal realist is in a position to state principles of plenitude in purely non-modal terms. Divers and Melia offer the following non-modal formulation of the principle of recombination, for example:

> For any individuals x_1, x_2, ... x_n there is a world containing any number of duplicates of each, if there is a spacetime big enough to hold them all, and such that for any spatiotemporal relation the duplicates in question stand in that relation.[18] (2002: 16)

Ersatzist theories of modality, by contrast, are not able to provide equivalent principles without making use of modal notions. The ersatzist needs a modal formulation of recombination such as the one offered by Lewis: anything *can* exist with anything else. The modal realist might try to argue that this means that primitive modality must enter twice into the ersatzist theory: both in the ersatzist analysis of modality (as truth according to some *consistent* set of propositions) and in the ersatzist principle of plenitude (that anything can coexist with anything else).

This line of argument is unlikely to succeed. Once ersatzists have paid the price for primitive modality in their analysis, they are in a position to apply the analysis directly to the principle of recombination. In general there is no need for a non-modal statement of principles of plenitude, as long as the modality involved in them is amenable to the same treatment as other modal claims.[19] Unless we are given special reason to believe that a theory of modality cannot handle the modality involved in some candidate principle of plenitude, we ought to allow principles of plenitude to be formulated in modal terms. And so, as in the case of recombination, ersatzists can incorporate a principle of miscellaneity into their theory. Again, this postulate must be stated in modal terms; but this is to be expected, given that

[18] This formulation of recombination is contestable on several grounds. It shares with Lewis's formulation a focus on individuals, so it imposes no recombination requirement on intrinsic monadic properties, and it also imposes no requirement on which spacetimes are possible. But it will suffice to make my point.

[19] This conclusion is also defended by Cameron 2011.

ersatzists make no claim to a reductive theory of modality. Of course, Lewis famously used the possibility of alien perfectly natural properties to argue against ersatzism.[20] But *that* problem is not that ersatzists cannot make sense of an infinitude of perfectly natural properties; it is that they are unable to make sense of merely quiddistic differences between pairs of worlds other than the actual world.

The prospects of using Generalized Plenitude* to distinguish between different theories of modality thus appear poor. All of the theories of modality on the table can match the performance of modal realism with respect to Generalized Plenitude* just by incorporating appropriate principles of recombination and of miscellaneity.

It may be possible to generalize plenitude* even further:

> ***Ultra-Generalized Plenitude*:*** Any metaphysical theory should incorporate plausible general principles concerning the extension of its subject-matter, and these principles should provide an optimal balance between simplicity and strength.

It is plausible that Lewis acknowledged the force of some condition like Ultra-Generalized Plenitude* in metaphysical methodology. In his correspondence with van Inwagen, he compares the plenitude challenge facing modal realists with related challenges facing set theory (1983a) and facing the theory of properties (1983b). How are we to say that there are all and only the sets there are? How are we to say that there are all and only the properties there are? In this analogy, the individual principles of plenitude are analogous to set-theoretic or property-theoretic comprehension principles. For Lewis, all of these principles are chosen to maximise overall theoretical adequacy while minimizing arbitrariness.

8. Conclusions

I have now answered all three of the questions with which I began. In answer to the interpretive question, I have suggested that principles of recombination, of miscellaneity and of inter/extrapolation are jointly constitutive of Lewis's plenitude criterion, and that this criterion can be motivated by the requirement that a modal realist theory should include a

[20] Lewis 1986: 159–65.

plausible candidate set of laws of the plurality. This allows us to give a positive answer to the metaphysical question: Lewisian modal realism does satisfy plenitude, by incorporating a plausible (although not precisely stated) candidate set of laws of the plurality. It also provides a basis for answering the methodological question: plenitude, thus interpreted, matters to us because it arises from an application of standard criteria of theory-choice to a modal realist's total theory. This accords well with Lewisian methodological naturalism. The interpretation of plenitude proposed here thus makes sense of all of Lewis's claims about the criterion, and of his motivation for proposing it; it also provides a principled basis for extending the criterion beyond Lewisian modal realism.[21]

References

Armstrong, David M. 1983. *What Is a Law of Nature?* Cambridge: Cambridge University Press.

Beebee, Helen and A.R.J. Fisher (eds). 2020. *Philosophical Letters of David K. Lewis Volume 1: Causation, Modality, Ontology*. Oxford: Oxford University Press.

---

[21] A distant ancestor of this chapter was written as part of my BPhil thesis in 2007 and a descendent was incorporated into my DPhil thesis in 2011; my thanks to Cian Dorr, John Hawthorne, Alex Paseau, Oliver Pooley and Robbie Williams for early guidance. Subsequently the material was presented at seminars in Melbourne, Manchester, and Oxford; I am grateful to these audiences for feedback, and especially to John Bigelow, Rohan French, Toby Handfield, Mike Hicks, Lloyd Humberstone, Nick Jones, Martin Pickup, Graham Priest, Dave Ripley, Alex Roberts, Gonzalo Rodriguez-Pereyra, and Zach Weber. I have also benefited from discussing plenitude-related material with Mark Balaguer, Sam Baron, Justin Clarke-Doane, John Divers, Heather Demarest, Michaela McSweeney, Jeff Russell, Wolfgang Schwarz, Henry Taylor, Barbara Vetter and Jenn Wang, and from comments by anonymous reviewers. Special thanks to Anthony Fisher for pointing me towards the relevant parts of Lewis' correspondence, and to Anthony and Helen Beebee for all their work on that correspondence, on the Manchester conference, and on this volume. This work forms part of the project A Framework for Metaphysical Explanation in Physics (FraMEPhys), which received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 757295). Funding was also provided by the Australian Research Council (grant agreement no. DP180100105).

Bigelow, John, Brian Ellis and Caroline Lierse. 1992. The World as One of a Kind: Natural Necessity and Laws of Nature. *British Journal for the Philosophy of Science* 43(3): 371–88.

Bricker, Phillip. 1996. Isolation and Unification: The Realist Analysis of Possible Worlds. *Philosophical Studies* 84(2/3): 225–38.

Cameron, Ross P. 2007. Lewisian Realism: Methodology, Epistemology, and Circularity. *Synthèse*, 156(1): 143–59.

Cameron, Ross P. 2009. What's Metaphysical about Metaphysical Necessity? *Philosophy and Phenomenological Research* 79(1): 1–16.

Cameron, Ross P. 2012. Why Lewis's Analysis of Modality Succeeds in its Reductive Ambitions. *Philosophers' Imprint* 12(8).

Chihara, Charles S. 1998. *The Worlds of Possibility: Modal Realism and the Semantics of Modal Logic*. Oxford: Clarendon Press.

Divers, John. 1999. A Genuine Realist Theory of Advanced Modalizing. *Mind* 108(430): 217–39.

Divers, John and Joseph Melia. 2002. The Analytic Limit of Genuine Modal Realism. *Mind* 111(441): 15–36.

Dorr, Cian. MS. How to be a Modal Realist. Available at https://philpapers.org/rec/DORHTB.

Fine, Kit. 2002. The Varieties of Necessity. in *Conceivability and Possibility*, ed. Tamar Szabó Gendler. New York: Oxford University Press, 253-281.

Jackson, Frank, 1998, *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Clarendon Press.

Knobe, Joshua, Ken D. Olum and Alexander Vilenkin. 2006. Philosophical Implications of Inflationary Cosmology. *British Journal for the Philosophy of Science* 57(1): 47–67.

Lewis, David K. 1973. *Counterfactuals*. Oxford: Blackwell.

Lewis, David K. 1983a. Letter 161. To Peter van Inwagen, 16 February 1983, in Beebee & Fisher 2020: 312–15.

Lewis, David K. 1983b. Letter 163. To Peter van Inwagen, 16 April 1983, in Beebee & Fisher 2020: 318–19.

Lewis, David K. 1983c. Letter 164. To Allen P. Hazen, 19 April 1983 in Beebee & Fisher 2020: 319–20.

Lewis, David K. 1983d. Letter 167. To Peter van Inwagen, 23 May 1983, in Beebee & Fisher 2020: 327.

Lewis, David K. 1986. *On the Plurality of Worlds*. Oxford: Blackwell.

Lewis, David K. 1987. Letter 183. To Jaegwon Kim, 31 March 1987, in Beebee & Fisher 2020: 356–9.

Lewis, David K. 1996. Maudlin and Modal Mystery. *Australasian Journal of Philosophy* 74(4): 683–4.

Lewis, David K. 2009). Ramseyan Humility. in *Conceptual Analysis and Philosophical Naturalism,* eds. David Braddon-Mitchell & Robert Nola. Cambridge:MIT Press, 203–22.

Maudlin, Tim. 1996. On the Impossibility of David Lewis' Modal Realism. *Australasian Journal of Philosophy* 74(4): 669–82.

Parsons, Josh.MS. 'Against Advanced Modalizing'. Available at: https://www.academia.edu/1016958/Against_advanced_modalizing.

Paseau, Alexander.2006. Genuine Modal Realism and Completeness. *Mind* 115(459): 721–9.

Price, Huw and Frank Jackson. 1997. Naturalism and the Fate of the M-Worlds. *Proceedings of the Aristotelian Society: Supplementary Volume 71*: 247–82.

Sider, Theodore. 2003. Reductive Theories of Modality, in *The Oxford Handbook of Metaphysics,* ed. Michael J. Loux. Oxford: Oxford University Press, 180-208.

Sider, Theodore. 2011. *Writing the Book of the World*. Oxford: Oxford University Press.

Van Inwagen, Peter. 1980. Indexicality and Actuality. *Philosophical Review* 89(3): 403–26.

Wilson, Alastair. 2020. *The Nature of Contingency: Quantum Physics as Modal Realism*. Oxford: Oxford University Press.