

Lewis's Empiricism*

Wolfgang Schwarz

Draft, 28 June 2019

1 Introduction

Lewis's epistemology, along with his philosophy of mind and philosophy of language, leaves no room for enquiry into non-contingent matters. Yet much of Lewis's metaphysics seems to be engaged in precisely this kind of enquiry. What did Lewis think he was doing? And whatever he thought, how can we make sense of what he was doing? More broadly still, how should we think of metaphysical enquiry if we like the kind of epistemology Lewis defended?

These are the questions I want to tackle. I'll begin by reviewing Lewis's epistemology. It is in many ways an empiricist epistemology, of the kind that made other philosophers skeptical or hostile towards metaphysics. Yet Lewis was neither skeptical nor hostile towards metaphysics. I'll discuss two ways of resolving this tension (without abandoning the empiricist epistemology); then I'll talk about what difference the resulting picture makes for Lewis's philosophy.

2 The empiricist epistemology

Lewis's epistemology starts with a simple and attractive idea. When we're ignorant of something, we don't know whether the world around us is one way rather than another. So we can characterise an agent's state of knowledge or belief by specifying which possibilities are compatible and which are incompatible with the agent's epistemic or doxastic state, without invoking mysterious entities somehow corresponding to *that*-clauses in ordinary-language attitude reports.

This attractive idea can be developed in several ways. In Bayesian models, which Lewis preferred, belief states are represented by a probability measure over the set of all worlds (for Lewis: over the set of maximal fusions of spatiotemporally related individuals). These states change primarily by conditionalisation: when new evidence arrives from the senses, the agent becomes certain of the evidence and renormalises the other probabilities.

In [Lewis 1996], Lewis also outlines a non-Bayesian epistemology. Here the content of an agent's knowledge is identified with the set of worlds in which the agent has the same

* This paper is a slightly improved version of a talk I gave at the *David Lewis and His Place in the History of Analytic Philosophy* Conference in Manchester on 19 June 2019. I should add some more citations and qualifications.

evidence as they actually have, minus worlds that are “properly ignored” in the context of attribution.

Both accounts are committed to what Richard Jeffrey calls “hardcore empiricism”: there are assumed to be special *evidence propositions* given by experience, known with absolute certainty, and providing the basis for all our knowledge. In the non-Bayesian model of [Lewis 1996], my current evidence proposition is the set of worlds in which I have the same perceptions and (quasi-)memories as in the actual world. This is a proposition I know even under intense skeptical pressure, when no possibilities are ignored.

Given that perceptions and (quasi-)memories are brain states, Lewis’s position in [Lewis 1996] seems to imply that we all know what brain states we’re in. But we don’t. So my current evidence can’t be the set of worlds in which I have my actual perceptions and (quasi-)memories, if these are types of brain states. It’s strange that Lewis didn’t see this problem. Anyway, I won’t dwell on it here. In the Bayesian model, one simple answer is to trade classical conditionalisation for Jeffrey conditionalisation, which does not assume that acquiring information makes anything certain. (Lewis expressed some sympathy for this move in [Lewis 1986b], but only for non-ideal creatures.) The result would be a version of Jeffrey’s “softcore empiricism”.

Jeffrey doesn’t explain why he uses the label ‘empiricism’. But here’s an empiricist element of all the models I have reviewed: genuine epistemic progress always comes from the senses. There’s no provision for *reasoning* to change subjective probabilities, or to change what is known. If you learn something, that’s always because of new information coming in from outside.

There are many other ideas traditionally associated with the label ‘empiricism’. Not all of these are accepted by Lewis (or Jeffrey). For example, Lewis accepts that much of what we know and rationally believe goes beyond what’s logically entailed by our evidence. Whatever my current evidence proposition is, it doesn’t logically entail that the sun will rise tomorrow, or that I have hands. Nonetheless, Lewis thinks I may well know both of these things, and rationally assign them high credence. Lewis also has no reservations about positing unobservable entities, and doesn’t think all meaningful sentences should be definable in observational terms. He takes scientific theories at face value, and trusts the scientific method.

In the Bayesian model, these parts of his epistemology are encoded in rationality conditions for “ultimate priors”: an agent’s hypothetical probability measure before acquiring any information. (See the discussion of “inductive method” in [Lewis 1974].) In the non-Bayesian picture, similar standards are encoded in some rules for what’s “properly ignored”.

In either model, it follows that we have substantive a priori knowledge. For example, we may know a priori that the unobserved resembles the observed. Counter-inductive worlds should be given low prior probability, and they can normally be ignored. But

this is the kind of knowledge that goes away when you look at it (on the [Lewis 1996] account). It is also not conclusive knowledge: counter-inductive hypotheses plausibly should not get probability zero.

There is also conclusive a priori knowledge in Lewis's models, but (setting aside technical problems in measure theory), there is only one example: the set of all worlds. This set is known by everyone, irrespective and independent of their evidence, and irrespective of what is ignored; it always has probability 1.

So even though there is some room for a priori knowledge, there is no room for a priori *enquiry*. You can't *come to know* a proposition through a priori reasoning (except by drawing away attention from worlds that can be properly ignored, in the model of [Lewis 1996], which is hardly substantive progress).

What shall we make of this? The fact that in the kinds of models Lewis presented an agent automatically knows all logical consequences of what they know – and therefore knows all logical truths – is called “the problem of logical omniscience”. It is commonly regarded as a serious flaw.

I disagree. I *like* this part of the models; I think it gets something importantly right. Ignorance is lack of information; if you lack information, you always lack information about the contingent world around us – there is no other kind of information to lack. The idea that there is a further, non-contingent part of reality (Plato's heaven? the metaphysical structure of reality?) to which we gain access by reasoning strikes me as mysterious and unnecessary.

This is not the place for an extended defence of logical omniscience. But let me note two points.

First, it is not easy to remove logical omniscience from Lewis's models. You might think logical omniscience is like the absence of friction and air resistance implied by simplified models of mechanics. Nobody seriously believes that there is no friction and air resistance, but leaving out these factors often simplifies a model, and the gain in simplicity can be worth the loss in accuracy. Logical omniscience is different. If you try to change Lewis's models so as to allow for ignorance of non-contingent matters – for example, by adding “impossible worlds” – you are likely to turn the original models into something altogether different that has lost any connection to the simple and attractive idea that motivated the original models (see [Bjerring and Schwarz 2016]).

Moreover, it is not enough to change the formal models. We also need to say what it is for an agent – ideally, described in non-intentional terms – to be adequately represented by a model. According to Lewis, an agent has such-and-such beliefs and desires iff they are in the state that plays the functional role his Bayesian model attributes to these beliefs and desires. If we want to allow for uncertainty towards non-contingent matters, we have to find a functional profile associated with that uncertainty. As Stalnaker pointed out, this is far from trivial (see [Stalnaker 1984], [Stalnaker 1991], [Stalnaker 1999]).

The second point I want to make is that our logical non-omniscience is not a datum. Granted, I don't know if there are infinitely many twin primes. Nobody knows. It's an open problem. So here is a non-contingent hypothesis – the twin prime conjecture – of which we don't know whether it is true. But that is not a proof of our logical non-omniscience. As Lewis rightly insists, the models he puts forward are not meant to supply a simplistic semantics of ordinary-language attitude reports. It is not part of the proposal that '*S* believes that *p*' is true iff *p* is true at all worlds that are doxastically accessible for *S*. The connection between attitudes characterised in terms of possible worlds and ordinary attitude reports is “complicated and multifarious” [Lewis 1986a: 34] (see also [Lewis 1981]).

One kind of complication that Lewis often mentioned (and that will become important later) arises if an agent is ignorant of certain linguistic facts. Suppose we pass a shop that says 'Ophthalmologist'; you ask, 'are ophthalmologists eye doctors?'; I say, 'I don't know'. What is the object of our ignorance? What are the possibilities we can't rule out? Plausibly, what we don't know is whether the word 'ophthalmologist' (as used in the community with which we're engaging) means *eye doctor*. For all we know, it might have this meaning, or it might have a different meaning. That's a contingent matter. So the (non-contingent) complement of the attitude report does not express the (contingent) object of our ignorance.

In general, if you find a true report of uncertainty or ignorance in which the complement sentence expresses a necessary truth, you have only found a counterexample to logical omniscience if the attitude report can be taken at face value. On Lewis's view, it often can't.

3 The tension

Let's recap. Lewis's preferred models of knowledge and belief leave no room for enquiry into necessary truths. Real knowledge and ignorance is always knowledge and ignorance of the contingent world around us.

This is what I mean by “Lewis's empiricism”. As I said, I like this aspect of his philosophy. I suspect I like it more than Lewis himself.

The tension with Lewis's metaphysics should be easy to see. When Lewis does metaphysics, he mostly discusses non-contingent matters: whether there is an objective distinction between natural and unnatural properties, whether there are universals or tropes, whether the past is as real and concrete as the present, whether other possible worlds are as real and concrete as the actual world, and so on. Lewis treated these questions as substantive, and assumed that they should be addressed by methods similar to those by which we address scientific questions about the world around us.

According to Lewis, when we evaluate different answers to a metaphysical question –

say, about the nature of possible worlds – we have to balance the costs and benefits. All else equal, an answer is good to the extent that it is simple, systematic, and conservative: if it respects common sense and previous opinion. Philosophical arguments can reveal hidden costs, or hidden benefits, and thereby change the credibility of theories. But even at the end of enquiry, there will rarely be an uncontroversial, decisive winner. Opposing views can always be upheld. “Once the menu of well-worked out theories is before us, philosophy is a matter of opinion.” [Lewis 1983].

If the object of study is really non-contingent, Lewis’s empiricist epistemology would predict a very different kind of attitude. Any well-defined non-contingent hypothesis expresses either the empty set or the set of all worlds. Why should the choice between such hypotheses turn on the balancing of theoretic virtues? How could it be a matter of opinion?

More strikingly perhaps, Lewis often professes ignorance of metaphysical issues. In [Lewis 1986a], he says that he doesn’t know whether there are duplicate worlds, or how large possible spacetimes can be, or whether positive and negative charge are co-instantiated at some world. Lewis is not at all troubled by this. “[W]hy should I think that I ought to be able to make up my mind on every question about possible worlds, when it seems clear that I may have no way whatever of finding out the answers to other questions about noncontingent matters – for instance, about the infinite cardinals?” [Lewis 1973: 89].

Something must go. There are two options. One is to go rationalist and accept that besides knowledge and ignorance of contingent matters, there is also knowledge and ignorance of non-contingent matters – a kind of knowledge ignored by Lewis’s official models of knowledge and belief. For some reason, similar epistemic standards apply in both domains, even though epistemic progress in one domain requires perception while progress in the other domain merely requires an armchair.

This response is worth pursuing. It might have been Lewis’s view (see esp. [Lewis 1986a: 109–115]), but the exegetical evidence is inconclusive (and the line between this response and the one I’m going to defend can be hard to make out). Nonetheless, I want to resist it. I would rather hold on to the empiricist epistemology and find some other way to accommodate metaphysical enquiry.

4 The reinterpretation strategy

In section 2 I mentioned that the surface form of attitude reports can be misleading: from the truth of ‘*S* is ignorant of *p*’ we can’t infer that the embedded clause *p* is false at some world compatible with *S*’s knowledge. My ignorance of whether ophthalmologists are eye doctors might just be ignorance of how a certain word is used in my community. Can we apply this trick to explain away our apparent ignorance of metaphysical issues? On

this approach, the true objects of metaphysical enquiry would not be the non-contingent hypotheses Lewis ostensibly discusses, but contingent surrogate propositions, perhaps about words or concepts.

This interpretation looks more attractive for some parts of Lewis's metaphysics than for others. The parts for which it looks more attractive are the parts where Lewis offers a conceptual analysis of causation, laws of nature, chance, dispositions, knowledge, belief, conventions, values, and so on.

Lewis usually puts his analyses in the material mode: " C causes E iff so-and-so". On the present proposal, the true object of his proposal is a contingent hypothesis about our words or linguistic dispositions: something like the hypothesis that that we are disposed to apply the word 'cause' to a pair C, E iff so-and-so.

(The claim is not that this contingent hypothesis is semantically expressed by the original statement, ' C causes E iff so-and-so'. Rather, the claim is that asserting this statement somehow conveys a belief in the contingent proposition, perhaps through a process of "metalinguistic diagonalisation", as Stalnaker suggested.)

This reinterpretation makes sense of our practice of discovering and responding to counterexamples. Suppose at some point we believed that knowledge is justified true belief, and then we came across a Gettier case. Our original belief plausibly predicts that we should be disposed to classify all (relatively simple) cases of justified true belief as knowledge. That's why data about people's intuitions is relevant to philosophical analysis. (If the target hypothesis were a non-contingent hypothesis about the nature of knowledge, by contrast, it is unclear why contingent information about people's dispositions to assent to certain strings of symbols would be relevant.) When we consider the Gettier example, we notice that we are not at all disposed to classify the case as 'knowledge', even though we do judge it to be a case of justified true belief. We made a surprising discovery about our linguistic dispositions, and thereby plausibly about other people's dispositions, and thereby about the meaning of the word 'knowledge'.

The reinterpretation strategy also make superficial sense of Lewis's methodology.

For one, it's understandable why conservatism plays a role: if an analysis deviates widely from our prior opinions, it is unlikely to capture our dispositions for applying the relevant terms. As Lewis says in [Lewis 1997]: "It won't do to say that colours do not exist; or that we are unable to detect them; or that they never are properties of material things". If you say any of these things, it is doubtful you are using 'colour' in its normal sense.

Why should metaphysical enquiry follow similar standards as empirical enquiry: why should we favour simple, systematic, and elegant theories, why should we balance benefits and costs, etc.? The answer is easy. On the present interpretation, metaphysical enquiry *is* empirical enquiry. It's a special type of empirical enquiry one can do largely by reasoning and introspection.

Uncertainty and ignorance of metaphysical issues also makes perfect sense, and we can explain why there are no knockdown arguments, why contrary views can always be upheld. This is the well-known underdetermination of theory by observation. Faced with an apparent counterexample, you can always claim that our intuitions in this case are mistaken: that our concept of knowledge really does apply to Gettier cases, but that we get confused when we think about these cases. (We know that we are not perfect in applying our concepts. Think of judgements about whether a number is prime, or whether some statistical evidence supports some conclusion.)

There is some evidence that Lewis endorsed the reinterpretation strategy. In a letter to Mary Kate McGowan (published on <http://www.projects.socialsciences.manchester.ac.uk/lewis/letter-month-february-2018/>), he explains that conceptual analysis is an investigation into our classificatory dispositions. The goal, he says, is to uncover his dispositions, by developing “fallible hypotheses about an independent (mental) reality that I can’t examine directly”. Moreover, he often approvingly cites Stalnaker’s attempts to explain away apparent counterexamples to logical omniscience in terms of contingent surrogate propositions. [John Bigelow reports that in conversation, Lewis also applied and defended the reinterpretation strategy for apparent cases of non-contingent knowledge.]

5 Problems with the reinterpretation strategy

According to the reinterpretation strategy, the true objects of metaphysical enquiry aren’t non-contingent hypotheses about metaphysical reality, but contingent hypotheses about our words or concepts. I don’t think this strategy fully succeeds at reconciling Lewis’s empiricist epistemology with his rationalist metaphysics. I’ll mention three problems.

First, let’s have a closer look at methodological virtues. If the true subject matter are our dispositions to use a word or concept, we should look for simple and systematic theories of *our dispositions*. This is not tantamount to looking for simple and systematic *analyses*. For example, a systematic hypothesis about why we are disposed to make certain judgements might appeal to universal cognitive mechanisms, such as our tendency to use spatial concepts when reasoning about non-spatial issues. Yet considerations like these are alien to Lewis’s metaphysics.

An especially puzzling virtue, from the perspective of the reinterpretation strategy, is ontological parsimony. Lewis thinks an analysis that posits fewer types of things is, all else equal, better (see e.g. [Lewis and Lewis 1970]). But why should ontological parsimony be relevant if the object of enquiry are our linguistic dispositions? Note that the ontological parsimony of an analysis does not make the surrogate proposition about our dispositions any more parsimonious.

A second, exegetical, problem is that Lewis seemed to treat some hypotheses in metaphysics and maths as not just unknown, but unknowable. (Remember the assertion about

large cardinals.) That's odd if the relevant hypotheses are really hypotheses about our linguistic dispositions. It is tempting to think that there's a recursive enumeration of all our linguistic dispositions, since they are generated by an effectively finite computational machine.

Third (and relatedly), there is more to metaphysics (and maths) than conceptual analysis. Take modal realism. Lewis says that other possible worlds are as real and concrete as the actual world. That doesn't look like an analysis of anything. What is the contingent surrogate proposition supposed to be?

It's not hard to understand the project Lewis is engaged in when he defends modal realism. The project starts with the variety of statements we would like to accept. These statements (sentences) appear to quantify over tables, persons, events, possible worlds, numbers, sets, and many other things. As a metaphysician (and Quine student), Lewis wants to know what we should make of these statements. Can apparent talk about possible worlds be paraphrased away, perhaps in a fictionalist manner? Lewis says no: only with serious costs. But it is useful talk, so we also don't want to give it up. So we must accept that there really are possible worlds. We should then ask what kinds of things they are. Are they *sui generis* entities that cannot be described in any other way? That's "magical ersatzism", and Lewis argues that it, too, has serious costs. Lewis suggests that our quantification over worlds is best understood as quantification over maximal spatiotemporally related mereological fusions.

In general, much of Lewis's metaphysics tries to systematise some part of the statements we are inclined to accept. He wants to see which terms can be defined by others, and how the things we talk about in one fragment of the language relate to the things we talk about in other fragments. The goal is to find a simple, elegant, perspicuous translation, with few primitives. In Lewis's final metaphysics, all first-order quantifiers ultimately range over constituents of possible worlds and mereological fusions of such constituents; the primitive concepts ("ideology") involve the machinery of classical monadic second-order logic, as well as a few concepts for parthood, fundamentality, identity, etc.

What's being systematised in this project are not propositions – sets of worlds. The starting point are sentences. The task is similar to the familiar mathematical project of axiomatising a given (mathematical or physical) theory, of finding a minimal set of primitive terms and axioms from which the entire theory can be generated by means of definitions and logical inference.

It's plausible that this project involves language. But I don't find it plausible that the project can be redescribed as one of finding out contingent facts *about* language.

The worry I have is perhaps easier to illustrate with a mathematical example. Consider a simple arithmetical task – say, to check if 29 is prime. To do this, you might check whether 29 is divisible by 2, then whether it's divisible by 3, then whether it's divisible by 5, and then note that $7^2 > 29$, so there's no point trying prime numbers from 7 onwards.

All these steps are easy. For example, to see that 29 is not divisible by 2, you only need to notice that the last digit, ‘9’, is not even. What does the reinterpretation strategy say about this little computation? At each step, it says, you acquire contingent information about the world. What is that information?

Well, whenever you figure out a mathematical truth S , you thereby also figure out that you’re disposed to assent to S (when given sufficient time to think). So we can indeed discover contingent facts about the world through pure reasoning.

Fair enough. But is this a plausible account of your epistemic progress? Does your progress as you try to figure out whether 29 is prime really consist in discovering facts about your dispositions, or about words? More pressingly, didn’t you already *have* the information you supposedly acquire? Couldn’t you all along rule out worlds in which the last digit of ‘29’ is not ‘9’ (if there are such worlds), and worlds in which you’re not disposed to say that a number that ends in ‘9’ isn’t divisible by 2? You may not have consciously thought of these facts before starting the computation, but it is hard to believe that you didn’t know them, in the semi-technical sense of the formal models reviewed above. (It is also hard to believe that you knew these facts, but in different “fragments” of your belief system. Since computations can get longer without bounds, there would have to be an unbounded number of fragments.)

Here’s what I think is missing in the empiricist models of knowledge and belief.

Physical agents must store information in a physical format. The format in which information is stored makes a difference to its use. Some representational formats make it easy to apply a piece of knowledge to a given task, others require lengthy computations. If a lengthy computation is needed, and you don’t have much time, you can’t perform the task, even though you have the information. You have it, but in the wrong form.

When asked if 29 (or 2529) is prime, I have all the information needed to answer the question. I don’t need to observe my dispositions or inspect Plato’s heaven to answer the question. But I need to do some computations to convert the stored information into a ‘yes’ or ‘no’ answer.

When we say that someone is ignorant of p , I think, we sometimes mean that they are not capable of performing a certain task that an ideal agent could perform on the basis of their knowledge of p . Often the relevant task is linguistic: assert or assent to ‘ p ’ (or a relevant synonym). At other times the task may be to make a certain move on the chess board, or to choose a certain path in a maze.

On this picture, much of what looks like a priori enquiry is manipulation of representational vehicles: we try to bring what we already know into a desirable form that allows us to perform a certain task.

The proposal is fruitfully combined with the reinterpretation strategy. Perhaps the true object of enquiry, when I wonder whether 29 is prime, is a contingent hypothesis about the words ‘29’ and ‘prime’. I already know that this hypothesis is true (in the

technical empiricist sense), since I have all the relevant information. But I don't have it in the right form. So I need to do some work to convert however the information is stored (plausibly, in the form of recursive rules, which are themselves stored in some lower-level manner) into an explicit answer to the question.

I'm not sure if the proposal also makes sense without the reinterpretation strategy. Could we say that the true object of mathematical or metaphysical enquiry is the set of all worlds, and the enquiry consists in manipulating representations of that set? Maybe.

In any case, much of metaphysics seems to fit this picture reasonably well. As I mentioned above, when we wonder whether there are universals or possible worlds, the project looks much like that of axiomatising a theory. We have a goal, perhaps inherited from science: to find a succinct, systematic (linguistic) description of reality, with few unexplained primitives in the ideology and a qualitatively sparse ontology. The task is to take the sentences we believe in and convert them into such a description.

Why do we seek a succinct description of reality? In empirical domains, simplicity is a fallible indicator of truth. Not so in the domain of the non-contingent, where there is no further goal of finding out the truth. (There's no truth to be found, except for the set of all worlds, and we've already found that.)

So it's a bit mysterious why similar theoretical virtues should apply to a priori and a posteriori enquiry. But it's not completely mysterious. Simplicity helps clear thinking. It's hard to reason in an intensional language with 257 primitives; it's easy to make mistakes in such a language, and to miss connections between the various primitives. And there's at least an indirect connection to truth: if you make a mistake, you are likely to accept sentences which express the necessary false proposition.

6 What difference does it make?

I have offered a suggestion for how to make sense of metaphysical (and mathematical) enquiry in an empiricist epistemology of the kind Lewis officially promoted.

As a historical point, I think it is clear that Lewis did not begin with any such conception of metaphysics (or maths).

Once again, modal realism provides a good illustration. Lewis began his career as a logician and philosopher of language. He noticed the pervasive parallels between quantification over worlds and quantification over time. In the case of time, he was attracted to the eternalist view on which other times are as real and concrete as the present. His early modal realism simply adapts that view to the modal: other worlds are as real and concrete as the actual world; 'actual' is an indexical like 'now'; truths about what could have happened are grounded in amodal truths about other worlds, just as truths about what did or will happen are grounded in tenseless truths about what happens at other times.

We cannot gain full and certain knowledge about other times just by inspecting the present. Similarly, Lewis seemed to think, we cannot gain full and certain knowledge of other worlds by inspecting the actual world. Thus his professed ignorance about duplicate worlds, the size of possible spacetimes, etc.

In *Plurality*, this is turned into one of Lewis's main arguments against ersatzism: the ersatzist who identifies possible worlds with (say) sets of sentences has to say which of these sets represent genuine possibilities and which do not. But that, Lewis argues, cannot be done without invoking primitive modality.

So Lewis accepted a kind of realism about modality that's logically independent of the ontological doctrine he called 'modal realism'. According to this kind of realism – let's call it *modal inflationism* –, the extent of logical space is a mind-independent, objective matter to which we have only imperfect epistemic access.

This modal inflationism does not fit my empiricist (re)interpretation of a priori metaphysics. The early Lewis failed to appreciate the crucial disanalogy between time and modality: what happens at other times is contingent, what happens at other worlds is not.

Lewis's inflationism also clashes with other aspects of his philosophy. For example, Lewis says that the possible worlds are a "philosopher's paradise", not just because they allow for a reduction of opinions about metaphysical modality. A more important role for possible worlds lies in the modelling of language and attitudes. (There's no analogous role for other times.) If there are no worlds with talking donkeys, then 'there are no talking donkeys' comes out analytic and a priori. But the sentence is obviously not analytic, nor does it express an a priori truth. So we can be sure that (if there are possible worlds at all then) there are worlds with talking donkeys. This is entailed by the role that defines the concept of possible world.

What about duplicate worlds? Let's assume that the job descriptions for possible worlds leave open whether there are duplicate worlds. Lewis seemed to think that there is nonetheless a fact of the matter, even though we can never find out what it is – not even in principle, abstracting away from our cognitive limitations. That's also what he seemed to think about large cardinals, and about the Continuum Hypothesis.

On the picture I've tried to paint, these metaphysical or mathematical conjectures should not be understood as conjectures about a mind-independent, non-contingent part of reality. The project of modal metaphysics is one of systematising a certain description of reality. If it turns out that the best systematisation does not settle the question of duplicate worlds, or if it turns out that there are equally good systematisations giving different answers, objective reality won't select one answer as right and the other as wrong. If there's a genuine tie, then there's no fact of the matter.

In this way, my picture of a priori enquiry supports a kind of relativism that I don't think Lewis would have liked. To take a more vivid example, suppose there is a rough tie

between postulating universals and postulating a primitive resemblance relation, in terms of whatever theoretical virtues govern metaphysical enquiry. Plausibly, these virtues are somewhat vague. Suppose you prefer the universals and I prefer primitive resemblance. Who is right? On my picture, there's no objective, "external" answer.

It is implausible that either answer is analytic by the lights of our standing linguistic conventions. When we systematise our descriptions of reality, we also tidy up our starting point. If you and I favour different ways of tidying up, nothing more can be said. It's not like one of us (or both of us) get the non-contingent metaphysical structure of reality wrong: there is no such thing as the non-contingent metaphysical structure of reality.

(But what if our tidied-up system postulates a non-contingent metaphysical structure of reality? For example, nothing I've said shows that Lewis's modal realism is not part of the best metaphysical system. If we *accept* modal realism – where "acceptance" is some kind of attitude towards sentences, not sets of worlds – we will accept various statements about spatiotemporally maximal fusions of individuals. And we will accept these statements as non-contingent. So we'll end up talking as if there's a non-contingent metaphysical structure of reality. And that's fine, I think. But we should be able to take a step back and get clear about the nature of the project, which is not usefully understood as uncovering hidden facts about metaphysical structure.)

References

- Jens Christian Bjerring and Wolfgang Schwarz [2016]: "Granularity Problems". *The Philosophical Quarterly*, 67(266): 22–37
- David Lewis [1973]: *Counterfactuals*. Oxford: Blackwell
- [1974]: "Radical Interpretation". *Synthese*, 23: 331–344
- [1981]: "What Puzzling Pierre Does Not Believe". *Australasian Journal of Philosophy*, 59: 283–289. Reprinted in [Lewis 1999]
- [1983]: *Philosophical Papers I*. New York, Oxford: Oxford University Press
- [1986a]: *On the Plurality of Worlds*. Malden (Mass.): Blackwell
- [1986b]: "Probabilities of Conditionals and Conditional Probabilities II". *The Philosophical Review*, 95: 581–589
- [1996]: "Elusive Knowledge". *Australasian Journal of Philosophy*, 74: 549–567
- [1997]: "Naming the Colours". *Australasian Journal of Philosophy*, 75: 325–342. Reprinted in [Lewis 1999]
- [1999]: *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press

David Lewis and Stephanie Lewis [1970]: “Holes”. *Australasian Journal of Philosophy*, 48: 206–212

Robert Stalnaker [1984]: *Inquiry*. Cambridge (Mass.): MIT Press

— [1991]: “The Problem of Logical Omniscience I”. *Synthese*, 89

— [1999]: “The Problem of Logical Omniscience II”. In *Context and Content*, 255–273