On the Probability of Plenitude

Jeffrey Sanford Russell

May 2020 draft

Abstract

I examine what the mathematical theory of random structures can teach us about the probability of *Plenitude*, a thesis closely related to David Lewis's modal realism. Given some natural assumptions, Plenitude is reasonably probable a priori, but in principle it can be (and plausibly it has been) empirically disconfirmed—not by any *general qualitative* evidence, but rather by our *de re* evidence.

> There is one thing needful: Everything.
> The rest is vanity of vanities.
>
> G.K. Chesterton

1 Incredulous Stares and Incredible Hypotheses

David Lewis (1986) held that reality is wildly vast and varied. Famously, this view prompted incredulous stares. Here's how Lewis put it:

> When modal realism tells you—as it does—that there are uncountable infinities of donkeys and protons and puddles and stars, and of planets very like Earth, and of cities very like Melbourne, and of people very like yourself, … small wonder if you are reluctant to believe it. And if entry into philosophers' paradise requires that you do believe it, small wonder if you find the price too high. (1986, 133, original ellipsis)

Lewis called his view *modal realism* because he also held a second thesis: a distinctive brand of *modal reductionism*. Eliding some complications, he held that what is

Thanks to Andrew Bacon, Cian Dorr, Maegan Fairchild, Jeremy Goodman, John Hawthorne, Yoaav Isaacs, Shieva Kleinschmidt, Alison Peterman, and an anonymous referee for helpful comments.

metaphysically possible just is what is true within one of these many worlds (1986, sec. 1.2). But this essay is not especially concerned with this part of Lewis's world-view. My present focus is just on what things are like, rather than what things *could* be like. Let's call the non-modal aspect of Lewis's view described in the quotation *Plenitude*. Very roughly, Plenitude is the view that there are very many things of very many kinds: talking donkeys and blue swans and nearly indiscernible counterparts of Earth and Melbourne and ourselves and many other things besides. In Section 2 I'll return to the issue of how to make Plenitude more precise.

Lewis acknowledges that Plenitude "does disagree, to an extreme extent, with firm common sense opinion about what there is" (1986, 133). Insofar as the deliverances of firm common sense opinion are rational, this is for one of two broad kinds of reason: a posteriori and a priori. (Insofar as the deliverances of firm common sense opinion are *not* rational, it strikes me as no serious cost to depart from them.) Perhaps Plenitude is implausible because it is empirically disconfirmed, either by ordinary experience or scientific investigations. I will return to this idea in Section 6; but note at the outset that it is not immediately obvious what these observations might be that would disconfirm Plenitude. Sections 2 and 3 will explore the alternative: that Plenitude is intrinsically implausible.

I will take it that there is such a distinction between hypotheses that are more or less credible, prior to experience, and that we can model this using the tools of probability theory. The starting idea is that there is some probability measure over propositions which represents their a priori probability—a "rational ur-prior."[1] (Perhaps there are many a priori rational probability measures, or perhaps what probabilities are rational for one person or community differs from what is rational for another, or perhaps it is a vague matter which of many different probability measures is rational. Let's ignore those complications for now.)

Some striking mathematical results from the theory of *random structures* seem to bear directly on the probability of Plenitude. A theorem from Paul Erdös and Alfréd Rényi says that *almost all* structures satisfy Plenitude: read naïvely, this theorem seems to say that the probability of Plenitude is maximally high![2] What is more, the same theorem shows that almost all structures are *isomorphic* to one another; so it also seems that the probability of *everything* interesting must be one or zero. (Quite a coup for rationalism!) It is understandable that some mathematicians have called this theorem "paradoxical" (Cameron 1997).[3]

[1] This is broadly in the spirit of Carnap (1950). See Meacham (2016) and references therein.

[2] Erdős (1959); Erdős and Rényi (1963). For a model-theoretic perspective see Hodges (1997, sec. 6.1); for a computational perspective see Libkin (2013 ch. 12).

[3] I use the word "structure" in the sense from model theory (see Section 2)—these are also called "models." Two numerically distinct structures can be "structurally exactly alike"—that is, isomorphic.

Of course, we should be very cautious about drawing bold epistemological conclusions from pure mathematics. My goal is to carefully examine these random structure results in order to make their assumptions clear and learn what they have to teach us about the credibility of Plenitude. The "paradoxical" implications drawn from a naïve reading of the mathematics must be moderated: in fact, I do not think that Plenitude has probability one. But even when appropriately qualified, the math teaches us some interesting things. Natural-seeming "Humean" assumptions push us toward the view that Plenitude is not so improbable *a priori* (Section 3). But our investigation also illuminates how empirical evidence might tell against Plenitude, and thus justify incredulity *a posteriori* (Section 6). Along the way (Sections 4 and 5), I'll also present some interesting generalizations and consequences of the standard results.

While Plenitude may be at odds with common sense opinion, it has many voices in its favor throughout the history of philosophical and scientific opinion. Lovejoy (1936, 52) introduced the term "the principle of plenitude" for the intertwined ideas

> that no genuine potentiality of being can remain unfulfilled, that the extent and abundance of creation must be as great as the possibility of existence and commensurate with the productive capacity of a 'perfect' and inexhaustible Source, and that the world is better the more things it contains.

He traces this influential cluster of ideas from Platonist roots through neo-Platonism, scholasticism, and rationalism. Plenitude is theologically loaded: it was deduced as a consequence of the goodness of God, and it also played a role in theodicy.[4] Contemporary physicists have joined historic theologians, in speculations about the "multiverses" of Everettian quantum mechanics or inflationary cosmology.[5] While I hope that the investigation I am taking up will indirectly illuminate these other ideas, I won't be exploring the connections to theology or physics here.

2 Plenitude and Patterns

Our first task is to formulate Plenitude a bit more precisely. Part of the Lewisian worldview is that there are certain *perfectly natural*, or *fundamental*, properties and re-

[4]For one example, Thomas Aquinas (following Plotinus and Augustine) writes that "the perfection of the universe is attained essentially in proportion to the diversity of natures in it, whereby the divers grades of goodness are filled …" (*Commentary on the Sentences*, book I, dist. 44, as quoted in Lovejoy 1936, 77). Compare more recent "multiverse theodicies" (McHarry 1978; Kraay 2010 inter alia).

[5]See for example Wallace (2012) and Knobe, Olum, and Vilenkin (2006). The "mathematical universe hypothesis" of Tegmark (2008) has an especially close family resemblance to the Plenitude thesis I'll explore here.

lations. (We won't fuss about what these might be.) I'll be investigating variations of this thesis:

Pattern Plenitude. Every pattern of fundamental properties and relations is instantiated.

We could arrive at Pattern Plenitude by Lewis's route, putting together two ideas. First, the *combinatorialist* idea that every pattern of fundamental properties and relations *could* be instantiated (Armstrong 1989; Lewis 2009; see also Saucedo 2011; for critical discussion see Wang 2013). Second, the *modal realist* idea that if a pattern *could* be instantiated, then it *is* instantiated (within some world). But note that while it can be deduced from modal premises, Pattern Plenitude is not itself a modal principle.

To make the idea of a "pattern" more precise, I'll follow Russell and Hawthorne (2018) in using relational structures. For technical simplicity we will restrict our attention to some finite set L of fundamental relations. (I'll count properties as one-place relations.) An **L-structure** S consists of a set of objects D—the **S-domain**—and a function that takes each n-place relation F in L to a set of n-tuples of elements of D—the **S-extension of F**. Structures are abstract mathematical objects; but some special structures correspond to genuine patterns of properties and relations in the world. We say that an L-structure S is **instantiated** iff there is some set D' such that S is *isomorphic* to the pattern of L-relations which is really displayed by the individuals in D'.[6]

Here is our refined principle:

L-Pattern Plenitude. Every L-structure is instantiated.

This version of Plenitude is a bit different from the way Lewis put things. Lewis originally presented his own theory of Plenitude in terms of a *recombination principle*, which Daniel Nolan concisely summarizes thus: "for any objects in any worlds, there exists a world that contains any number of duplicates of all of those objects … size and shape permitting" (Lewis 1986, 89; Nolan 1996, 239). The relationship between Pattern Plenitude and Lewis's "cut and paste" style recombination principle is delicate (see Wilson 2010; Russell and Hawthorne 2018). But let's not get sidetracked by these details. Though there are differences in detail between Pattern Plenitude

[6]That is, there is some one-to-one correspondence f from D' to the domain of S such that, for each n-place relation F in L and any individuals $d_1, \ldots, d_n$ in D', $d_1, \ldots, d_n$ stand in F iff $(fd_1, \ldots, fd_n)$ is in the S-extension of F.

(explicated in terms of structures) and Lewisian recombination (explicated in terms of duplication), they have at least this much in common: they are views that say that reality is vast and varied, that there are blue swans and talking donkeys and alternative Melbourne-counterparts and so on.[7] Thus defenders of Pattern Plenitude are just as liable to be subjected to incredulous stares.

Lewis added a caveat to his recombination principle: "size and shape permitting" (1986, 89). This was in part to avoid cardinality paradoxes that arise for the unrestricted version (Forrest and Armstrong 1984; Nolan 1996; Russell and Hawthorne 2018). Some related issues arise for Pattern Plenitude. A structure can have any set as its domain. The domain could consist of people, or electrons, or numbers, or pure sets. (Remember that for such a structure to be *instantiated* does not require that pure sets themselves stand in the natural relations in L; rather, it just requires that certain n-tuples of pure sets are *isomorphic* to the pattern of relations that some other individuals bear.) But this means that unrestricted Pattern Plenitude requires that there are very, very many bearers of fundamental properties and relations: indeed, more than any cardinal number, too many to form a set. This explosion may not be disastrous,[8] but it might be more than we hoped for. It also raises technical difficulties—particularly when we try to reason about this hypothesis using standard tools from probability theory. So we'll begin our investigations with a more modest thesis, which (a bit like Lewis's recombination principle) incorporates a restriction on size.

Countable Pattern Plenitude. Every countable L-structure is instantiated.

That is to say, every pattern that involves at most as many objects as there are natural numbers is instantiated.

Countable Pattern Plenitude does not require any explosion in the size of concrete reality, but it is still a striking thesis. It too evidently has many of the consequences that inspired incredulity in Lewis's interlocutors; for I take it that blue swans, talking donkeys, Melbourne-counterparts, and so on don't essentially require uncountable infinities of fundamental-property-bearers.[9] Countable Pattern Plenitude is enough

[7] At least, Pattern Plenitude says this given that there is *some* way of configuring finitely many fundamental properties and relations which would amount to there being a blue swan, etc. This seems plausible, but it might be wrong.

[8] Russell and Hawthorne (2018, 177 and appendix B) show that unrestricted L-Pattern Plenitude is logically consistent (given a large cardinal axiom).

[9] But this also might be wrong. For example, maybe the only patterns that manifest blue swans involve relations among all of the continuum-many point-parts of an extended region of space-time. I am supposing that being a blue swan is compatible with having discrete micro-structure.

for there to be an awful lot of very strange things out there. It's true that we have elided one part of the claim that Lewis said inspired incredulity—that there are *uncountable infinities* of these strange creatures. But merely *countable* infinities of blue swans and Melbourne-counterparts also seems at least a bit wild.[10]

So now we have narrowed our investigation to this question: does Countable Pattern Plenitude have low a priori probability?

3 Humean Probabilities

Our next task is to work out a version of Erdös and Rényi's theorem that bears on this question. When mathematicians talk about what is "almost sure," they typically take some especially natural and symmetric probability measures for granted in the background. For philosophical purposes, though, we will want to be very cautious about this point—the probability measure is where the action is.[11]

Let's start with a simple analogy. Suppose you know that an angel has visited the squares on an infinite chess board one by one, and decided whether to color each one red or black by flipping a coin. Then you should expect each finite pattern of red and black squares to show up somewhere or other. Indeed, you should be almost sure of this, in the technical sense: for any finite pattern, the probability that it occurs somewhere or other on the infinite chessboard is one. Moreover not only is *each* finite pattern almost sure to occur, but it is almost sure that *every* finite pattern of red and black squares occurs.[12]

[10] I don't know how things will go for stronger, unrestricted principles of Plenitude. We face technical obstacles right from the get-go, since we will need a fancier sigma-algebra of propositions: the direct analogue of the "structure algebra" introduced in Section 3, the Borel algebra on the set of atomic propositions, does not contain any non-trivial qualitative propositions at all in the uncountable case, so it is not really suitable.

[11] For this reason, my technical approach is a bit non-standard. The standard approach to random structures uses the device of *asymptotic probability*: one considers, for each n, a probability measure μ_n on the (finite) set of size n structures (up to isomorphism); then one examines properties of the limit $\mu = \lim_{n\to\infty} \mu_n$. But I find it harder to draw epistemic lessons from this machinery, which involves not one probability measure on an infinite domain of possibilities, but rather an infinite set of probability measures on different finite domains of possibilities. It seems to be well-known folklore that the relevant features of asymptotic probability (such as the zero-one law discussed below) carry over to the coin flip measure on an algebra of countably infinite structures, but I've been doing a lot of reconstruction to figure out how the details go.

[12] Countable additivity implies that if propositions $A_1, A_2, \ldots$ are each almost sure, then so is their conjunction. (The conjunction is equivalent to the negation of the disjunction of countably many mutually exclusive propositions "n is the smallest number such that A_n is false," each of which has probability zero.)

There are two key ideas doing work in this example. The first idea is that the universe—in the example, the chessboard—is infinite. The second idea is that, a priori, the different finite parts of the universe are, as Hume puts it, "entirely loose and separate" ([1748] 2007, 7.2.27). How things are configured in one finite part of the universe does not settle what goes on in some other separate part. The main mathematical fact is a powerful extension of this simple observation.

First, some technical preliminaries. Let D be a countably infinite set. We will suppose that some a priori epistemic possibilities are represented by L-structures with domain D. We can identify *propositions*—the bearers of probability—with sets of such structures. We can pick out a suitable algebra of such propositions: the **structure algebra** $\mathscr{S}$. A **structure measure** is any (countably additive) probability measure defined on the structure algebra.[13]

We will use structure measures to model certain aspects of the epistemic lives of rational beings. But before we do that, let's clear up some points about what we are modeling. Propositions in the structure algebra represent epistemic possibilities; but not *every* epistemic possibility has a home in the structure algebra. First, the structure algebra does not represent possibilities where some D-things are absent. But whatever possible individuals D represents, it may well be epistemically possible for some of them not only to fail to bear fundamental properties and relations, but to fail to be anything at all. (For example, it might be epistemically possible that there *aren't* infinitely many individuals.) Second, all of the propositions in the structure algebra are restricted to a certain *subject matter.* All of them represent ways certain fundamental properties and relations might be distributed over certain possible individuals. But this is not to say we can have no rational uncertainty about other subject matters—what other individuals are like besides those represented by D, and also what things are like in other respects besides those represented by L.[14] A structure measure simply does not represent any probabilities for these other subject matters. In either respect, though, don't take the absence of representation to be a representation of absence. We are using a coarse-grained model to investigate some kinds of rational uncertainty; we are not thereby implying that these are the *only* kinds of rational uncertainty.[15]

[13]In particular, $\mathscr{S}$ is the smallest σ-algebra that includes all of the *atomic* propositions (as in Definition 1).

[14]Furthermore, the structure algebra only represents uncertainty about the fixed domain of individuals under a certain stock of *modes of presentation.* Under these special modes of presentation, identity and distinctness facts about the individuals in question are perfectly certain. But that isn't to say we might not be in doubt about these issues under other modes of presentation, which are not captured by structure measures.

[15]To be more precise, you can think of a structure measure as derived from a richer, more realistic probability measure in three steps. (1) Choose some possible individuals to be represented by D and

It is also worth pausing on a point about the metaphysics of this model. The objects in the domain D are meant to *represent* certain epistemically possible individuals—in particular, certain candidates for bearing fundamental properties, so you might want to think of them as something like fundamental particles. Propositions in the structure algebra are meant to represent the ways those possible individuals might be. Any old abstract objects are adequate for mathematically *representing* such possibilities. The natural numbers would be just fine—they don't have to really *be* particles to do this job. But as is usual with mathematical models, it is often convenient to speak loosely as if the domain of the mathematical structures really includes the individuals themselves.

Now let's get to work.

Hume writes, "All beings in the universe, consider'd in themselves, appear entirely loose and independent of each other" ([1739] 2007, III.i.i). We might regiment this idea by saying that distinct matters of fact are *probabilistically* independent a priori.[16][17] We might think of it like this: for each fundamental n-place relation F, an angel visits each n-tuple of individuals and decides whether they instantiate F by the flip of a fair coin. In this case, each "basic" fact has probability one-half, independently of any other "basic" facts. This presents us with the following simple model. (I carry on a grand tradition of multiplying ahistorical senses of the word "Humean.")

Definition 1. A proposition A is **atomic** iff there is some n-place relation F in L and there are some individuals $d_1, \ldots, d_n$ in D such that A is the set of all structures S in $\mathcal{S}$ such that $(d_1, \ldots, d_n)$ is in the S-extension of F. A **basic** proposition is either an atomic proposition or its negation.

some properties and relations to be represented by L (as well as appropriate modes of presentation). (2) *Conditionalize* on the hypothesis that each of the D-individuals is there. (3) *Marginalize* on the subalgebra of propositions that are entirely about the distribution of L-properties over D-individuals.

[16] It is interesting to note that the term "independent" was used with its modern technical meaning at the time of Hume's writing: de Moivre (1738, 5–6) writes,

> If the obtaining of any Sum requires the happening of several Events that are independent on each other, then the Value of the Expectation of that Sum is found by multiplying together the several Probabilities of happening, and again multiplying the product by the Value of the Sum expected. …
>
> Two Events are independent, when they have no connexion one with the other, and that the happening of one neither forwards nor obstructs the happening of the other.

[17] Compare Popper ([1934] 1959, 379): "Every other assumption would amount to postulating *ad hoc* a kind of after-effect; or in other words, to postulating that there is something like a causal connection between [two basic propositions]."

Definition 2. A structure measure P is **strongly Humean** iff for any logically independent basic propositions $A, B_1, \dots, B_n$,

$$P(A \mid B_1 \cdots B_n) = \frac{1}{2}$$

As it turns out, this condition uniquely characterizes one particular structure measure—the "coin flip measure"—and it is an especially natural choice. But while it is especially simple and natural, the strong Humean condition is much stronger than we really need for the results that follow. The angel's coin need not be fair, and we don't need to assume that learning about how certain relations are configured tells you *nothing* about other cases—just not too much. The key assumption is just that certain conditional probabilities are not too high or too low. Here is a more relaxed version of the Humean condition.

Definition 3. Let ε be a probability. A structure measure P is **ε-Humean** iff, for any logically independent basic propositions $A, B_1, \dots, B_n$,

$$\varepsilon < P(A \mid B_1 \cdots B_n) < 1 - \varepsilon$$

A structure measure P is **moderately Humean** (or just *Humean* for short) iff P is ε-Humean for some $\varepsilon > 0$.

I'll return to the question of whether this Humean condition is reasonable later; first, let's see what it can do.[18] In these terms, here is what Erdös and Rényi's theorem tells us. (A proof sketch is given in Appendix A.)

Theorem 1. There is a countable structure S, **the random L-structure**, such that

(a) S satisfies Countable Pattern Plenitude.

(b) Any moderately Humean structure measure assigns the set of structures that are isomorphic to S probability one.

[18] Note that the moderate Humean condition is stronger than yet another way you might explicate Hume's idea: namely, that there are no a priori *necessary* connections between "distinct existences"—this basically corresponds to being ε-Humean for $\varepsilon = 0$.

This has two immediate consequences.

Corollary 1. For any moderately Humean structure measure, Countable Pattern Plenitude has probability one.

In *almost all* countable structures, every countable pattern of fundamental relations shows up somewhere or other. Plenitude isn't an exotic possibility in the space of possible infinite structures: it is the norm.

The second consequence concerns *qualitative* (or *general*) propositions: roughly, those that do not make reference to any particular individuals.

Definition 4. A proposition A in the structure algebra is **qualitative** iff A is invariant under isomorphisms, in the sense that for any structure S in A, every structure isomorphic to S is also in A.

Corollary 2. If P is a moderately Humean structure measure, then for every qualitative proposition A, either $P(A) = 0$ or $P(A) = 1$.

Corollary 2 leaves the Humean without any uncertainty about general qualitative matters. But note that it does *not* leave them completely certain on *every* matter. There is still plenty of room for doubt regarding "haecceitistic" *de re* propositions about how particular things are configured. I will return to this important point in Section 6.

I find these mathematical facts very striking. The way they work is by precisely capturing and extending a thought which many people find intuitively natural: in an infinite universe, whatever might happen probably will, somewhere (compare, for example, White 2018). The theorem goes beyond this intuition. It teaches us that to get, *almost surely*, *all* countable structures (including, say, structures manifesting infinitely many talking donkeys) we just need countably many individuals whose properties are sufficiently independent. It also teaches us that a certain qualitative *kind* of plenitudinous structure is almost sure to be instantiated.

So far this is just math: certain measures have certain features. What does it tell us about epistemology? Let's warm up by imagining someone whose opinions are guided by moderately Humean prior probabilities—call him Humberto. Humberto is sure there is a countable infinity of things whose fundamental properties and relations are “loose and separate.” Corollary 1 tells us that Humberto assigns Countable Pattern Plenitude prior probability one. Far from finding this hypothesis to be intrinsically implausible, he rather finds it *maximally plausible*. The hypothesis that there are *not* infinitely many talking donkeys brings forth his incedulous stares.

How about Humberto’s *posterior* probabilities? If the prior probability of Plenitude is one, then given any evidence that has non-zero prior probability, the posterior probability of Plenitude is also one. In general, for any propositions H and E, if $P(H) = 1$ and $P(E) \neq 0$, then the conditional probability $P(H \mid E) = 1$ as well. Standard Bayesian conditionalization can never make the almost-sure less than almost-sure. Furthermore, moderately Humean probabilities give positive probability to any finite conjunction of *basic* propositions—any particular local matter of fact, so to speak. So no “local” evidence like *this* can disconfirm Plenitude for a Humean like Humberto. If Humberto only has local evidence, then his posterior probability for Plenitude is also one.

Still, maybe Humberto has relevant evidence with prior probability *zero*. That takes us off the map of standard Bayesian conditionalization and into the wilds of non-standard probability theory. (But some strong arguments point us this direction anyway: Hájek 2003.) In that case, Humberto’s prior probabilities would not determine his posterior probabilities: we would also have to explore Humberto’s primitive *conditional* probabilities (or something similar). What sort of probability-zero evidence might Humberto have? Someone with a sufficiently liberal view of what counts as evidence might argue that Humberto’s evidence can include negative existentials like *there are no talking donkeys*. One such liberal view is that “knowledge, and only knowledge, constitutes evidence” (Williamson 2000, 185). Note, though, that Humberto’s path to any such general knowledge could not go by way of incremental confirmation, adding probability little by little until it becomes sure. Gaining this kind of evidence would have to take him all the way from probability zero to probability one in one great leap.

Should we be like Humberto? I see no persuasive reason to think so, despite the pleasant symmetry of Humberto’s priors. There might well be only finitely many bearers of fundamental properties and relations; and even if the universe is infinite, we might well expect its parts to be linked up somehow (more on this in Section 6). But learning mathematical facts about Humean structure measures hasn’t been a waste of time. Even if it isn’t what the priors command, one could be moved by *evidence* to share Humberto’s Humean outlook. Our priors plausibly include a Humean part, so to speak, and the theorems teach us important things about this part.

Call H a *Humean hypothesis* iff *conditionalizing* the rational ur-priors on H yields a moderately Humean structure measure. Corollary 1 tells us that the conditional probability of Countable Pattern Plenitude, given any Humean hypothesis, is one. It follows by the probability calculus that the a priori probability of Countable Pattern Plenitude is at least as high as that of any Humean hypothesis, and indeed it is at least as high as the disjunction of all Humean hypotheses.

Furthermore, it seems that Humean hypotheses do not deserve especially low a priori probability: while it seems far from *certain* that there be a countable infinity of things whose fundamental properties and relations are loose and separate, it also seems far from *absurd*. (But I don't really know how to argue for this, if it doesn't strike you the same way.) If that's right, then despite its extravagances, Countable Pattern Plenitude is not intrinsically especially implausible. Accepting this does not put us in the same boat with Humberto, though. As long as Humean hypotheses have probability less than *one*, even very high prior probability does not stand in the way of Plenitude being disconfirmed by a posteriori evidence in the standard way. We will investigate how this can work in Section 6; first, though, let's look at a generalization (Section 4) and a further consequence (Section 5) of Humean probabilities.

4 A Priori Constraints

I have been considering how things go with relations which are "entirely loose and separate." But there may be a priori constraints on the structure of certain fundamental relations. For example, perhaps certain structural facts about space and time are a priori certain (for example Kant [1781] 1999, A24/B39). Maybe it's a priori certain that the relation *before* is transitive and irreflexive, or that *one meter apart* is symmetric. Maybe it's a priori certain that nothing has exactly one kilogram mass and also has exactly two kilograms mass.[19]

It turns out that there is a broad family of "nice" theories for which the main morals of the theory of random structures still apply, and which give rise to their own versions of Plenitude. For a theory T, let **Countable T-Plenitude** be the proposition that every countable model of T is instantiated: that is, every small enough pattern of L-relations which is compatible with the "external" constraint T shows up somewhere or other. It turns out that if T is a reasonably nice general theory (in a sense we can make precise) then T has a "random model": this is a countable model that embeds every countable model of T. Reasonably nice theories include, for example, that every electron has mass, that *before* is transitive and irreflexive, and that *one meter apart* is symmetric.

[19] Wang (2013) raises analogous ideas as challenges to combinatorialism about metaphysical possibility.

We can formulate the main "niceness" condition on T in two equivalent ways: model-theoretically, or syntactically. Very roughly and intuitively, the model-theoretic formulation says that any two ways T might obtain can both obtain together. This is called the *Amalgamation Property*, and it has a very similar flavor to Lewis's "cut and paste" recombination principle. Similarly roughly and intuitively, the syntactic formulation says that T can be expressed in a "fundamental language" with predicates for each L-relation by universal generalizations which are "not too disjunctive." Precise statements are given in Appendix A.

We can also generalize the notion of *Humean* probabilities.

Definition 5. Let T be a theory. A structure measure is **moderately Humean within** T iff $P(T) = 1$ and there is some $\varepsilon > 0$ such that, for any basic propositions $A, B_1, \ldots, B_n$, which are jointly consistent with T,

$$P(A \mid B_1 \cdots B_n) > \varepsilon$$

We can prove that, if T is a *nice* theory, then any structure measure P which is moderately Humean within T assigns probability one to the random model of T, up to isomorphism. Thus the P-probability of Countable T-Plenitude is one, and every qualitative proposition has P-probability zero or one. (See Theorem 3 in Appendix A.)

5 Incredible Isolation

There's more to Lewis's worldview than just Plenitude: Lewis also held that worlds are *isolated* from one another (1986, sec. 1.6). There is no hope of traveling from Los Angeles to one of its better-planned counterparts in another world, free of the downtown 101–110–10 freeway interchanges: for in fact, Los Angeles bears no spatio-temporal relation at all to its otherwordly counterparts. More generally, according to Lewis, parts of distinct possible worlds bear no "analogically spatio-temporal" relations to one another. Lewis also proposes (though does not quite endorse) the further claim that parts of distinct worlds are not linked by any perfectly natural relations at all (1986, 76).

What can we say about the intrinsic plausibility of this doctrine, that distinct parts of the pluriverse are relationally isolated? First, let's restate this using our present framework of relational structures. Let F be a binary relation in L. (Something similar goes for more-than-two-place relations.)

Definition 6. A structure S is **F-connected** iff every pair of objects is linked by some chain of objects pairwise related by F: that is, for each a and b in the domain of S, there are $d_1, \ldots, d_n$ in the domain such that $a = d_1$, $b = d_n$, and for each i, the pair (d_i, d_{i+1}) is in the S-extension of F.

Corollary 3. For any moderately Humean structure measure P, and for any binary relation F in L, the set of F-connected structures has P-probability one.[20]

Thus, while we have an argument that Countable Pattern Plenitude is intrinsically plausible, this argument does not suggest that *isolation* is plausible. The same kind of probabilistic combinatorial reasoning that makes it likely that every pattern of fundamental relations is instantiated somewhere or other, in an infinite domain, would also make it likely that every pair of individuals is *linked* by a chain of fundamental relations.[21]

We could try to recover isolation "by hand"—by considering priors with extra constraints as in Section 4. Here is a simple way of setting things up. Let one of the relations in L be a *worldmate* relation. Let the *World Theory* be the proposition that *being worldmates* is an equivalence relation, and that fundamental relations hold only among worldmates: for each relation F in L, for any individuals $x_1 \ldots x_n$ that instantiate F, each pair of x_i and x_j are worldmates. The World Theory turns out to be a "nice" hypothesis in the sense from Section 4. So the World Theory has a unique (up to isomorphism) "random model": call this the *World Model*. Priors which are Humean within the World Theory—that is, priors which treat distinct existences as loose and separate insofar as they are jointly compatible with the World Theory—assign the World Model probability one, up to isomorphism. Moreover, this World Model includes infinitely many isolated *worlds* (equivalence classes under the worldmate relation).

That sounds promising, but in fact it hasn't got us very far. The worlds of the World Model are not like the varied worlds of a Lewisian pluriverse—rather, each of these worlds is qualitatively indiscernible from the others. The World Model turns out to

[20] *Proof Sketch.* Theorem 1 is proved by showing that the *Extension Property* (Definition 9 in Appendix A) has probability one. Let S be a structure with the Extension Property and let a and b be elements of S. Then S has a substructure whose domain includes just a and b. Furthermore, this substructure can be extended by adding a third element c, and putting both (a, c) and (c, b) in the extension of F. By the Extension Property, the original $\{a, b\}$-substructure can be extended to another substructure of S which is isomorphic to this $\{a, b, c\}$-structure—which means that a and b are linked by an F-related chain of elements in S (indeed, a chain of length three).

[21] Some have argued on a variety of completely different grounds that Lewisian modal realism does not require isolated worlds, and perhaps does better without them (Yagisawa 1992; Noonan 2014; Dorr, n.d.).

consist of countably many disjoint isomorphic copies of the original random structure.[22] So while this model does give us a kind of isolation, this is only because in *addition* to the great interconnected plenitudinous structure in which we live and move and have our being, there are *also* infinitely many other qualitatively identical plenitudinous structures which are isolated from it.

That isn't to say that Lewisian isolation couldn't be probable a priori for some other reason. The point is just that, unlike Plenitude, this conclusion does not follow from any obvious "Humean" premise about priors.

6 Regularities and Induction

Lewis drew this moral from his recombination principle:

> Combinatorialism tells us that the laws of nature are contingent. Let it be a law that every F is a G; combinatorialism generates a possibility in which an F is not a G, so that this law is violated. (2009, 209, see also 1986, 91).

But Plenitude implies that the laws of nature are not merely *possibly* false—they are *false*. The unrestricted universal generalization "every F is a G" has exceptions—otherworldly exceptions, let us grant, but exceptions nonetheless. For a simple example, if we pretend that *raven* and *black* are fundamental properties, then Pattern Plenitude implies that not all ravens are black.[23]

More generally, Pattern Plenitude is incompatible with exceptionless general regularities. Some historic defenders of plenitude embraced this moral. For example, a central tenet of Margaret Cavendish's system is that "the parts of Nature are infinite, and have infinite actions"; this is an important basis for her general attack on universal scientific theories ([1666] 2001, 26, also p. 68; see also Peterman 2017, sec. 4.1).

We can put this general moral a bit more precisely. Consider a "fundamental language" with predicates for each fundamental relation in L, and names for each individual in D. We can straightforwardly talk about the probability of any sentence ϕ in this language, with respect to a structure measure, by associating ϕ with the set of structures S in which ϕ is true.

[22] That is, the random structure for the smaller signature that omits the worldmate relation.

[23] According to Lewis (1986, sec. 1.9), "laws" are regularities that only hold around here, in the spatiotemporally contiguous part of reality he dubs "the actual world"; the official deliverances of natural science use only *restricted* quantifiers. For critical discussion see deRosset (2011); Williamson (2013, xii).

Definition 7. A **simple generalization** is a sentence of the form

$$\forall x_1 \cdots \forall x_n \, \phi(x_1, \ldots, x_n)$$

where $\phi(x_1, \ldots, x_n)$ is a quantifier-free formula with predicates in L and without names. An **instance** of a simple generalization of this form is a sentence $\phi(d_1, \ldots, d_n)$ for any individuals $d_1, \ldots, d_n$.

No non-tautologous simple generalization is true in a plenitudinous structure. (This is because any non-tautologous quantifier-free L-formula $\phi(x_1, \ldots, x_n)$ is falsified by some finite L-structure—in particular, it is falsified by some substructure of the random L-structure.)

Corollary 4. Any moderately Humean structure measure assigns every simple generalization probability zero.

Consider Humberto again, who is guided by moderately Humean priors. Humberto regards every substantive empirical unrestricted generalization as a priori almost surely false. If Humberto goes around observing things, including lots of black ravens and no non-black ravens, you might expect that this would raise his confidence that all ravens are black. Not so. Humberto takes it to be almost surely false that all ravens are black. He also assigns the conjunction of any finitely many instances of a simple generalization positive prior probability. So his conditional probability for any substantive generalization, given any finitely many "local" observations, is still zero. For people like this, completely general empirical science doesn't even get off the ground. Humberto's "Humean" outlook has taken him from the idea that distinct existences are "loose and separate" to radical inductive skepticism (compare Popper [1934] 1959, 379–80; for critical discussion see Howson 1973).

In contrast to Humberto's predicament, it seems reasonable to hope that *generalizations are confirmed by their instances*. It would be nice to be able to gain evidence for the thesis that all ravens are black by observing that this is black-if-a-raven, and that is black-if-a-raven, and so on. Let's consider how this goes in the more realistic scenario where some, but not all, of logical space has a Humean flavor. Suppose that in addition to a *Humean hypothesis H*, we also have an *induction-friendly* hypothesis I, conditional on which some simple generalization can be confirmed by instances. Suppose we go around observing these instances. Then eventually the probability of the next instance, given those so far, will be *higher* conditional on the induction-friendly hypothesis I than it is conditional on the Humean hypothesis H. So these instances will provide relative confirmation for I rather than H. As we observe more and more instances, this relative confirmation will continue to accumulate: in fact, the probability of H tends to zero, and the probability of Plenitude dwindles with it. This holds whatever the prior probabilities of H and I may be, as long as they are not zero.

Proposition 1. Let P_H be moderately Humean, and let P_I be a structure measure which is *induction-friendly* in the sense that

$$\lim_{n\to\infty} P_I(G \mid E_1 \cdots E_n) = 1$$

for some simple generalization G and instances $E_1, E_2, \ldots$. For some probability $p < 1$, let

$$P = p \cdot P_H + (1-p) \cdot P_I$$

Then

$$\lim_{n\to\infty} P(G \mid E_1 \cdots E_n) = 1 \quad \text{and} \quad \lim_{n\to\infty} P(Q \mid E_1 \cdots E_n) = 0$$

where Q is the Countable Pattern Plenitude proposition.[24]

Even if we start out strongly inclined toward a Humean perspective—and thus toward Plenitude—as long as our priors also include an "inductive part," then by observing instances of a simple generalization, eventually Plenitude can be disconfirmed to an arbitrarily strong degree.

This might seem paradoxical. Plenitude implies that there are non-black ravens; but it also implies that there are lots of black ravens, and that there are arbitrarily huge regions (substructures) of the universe that are entirely free of non-black ravens. If Plenitude logically implies that such regions exist, how could observing one be evidence against Plenitude?

[24] *Proof.* Let E^n abbreviate $E_1 \cdots E_n$. By the odds form of Bayes' Theorem,

$$P(G \mid E^n) = q_n \cdot P_H(G \mid E^n) + (1-q_n) \cdot P_I(G \mid E^n) \qquad \text{where} \quad \frac{q_n}{1-q_n} = \frac{p}{1-p} \cdot \frac{P_H(E^n)}{P_I(E^n)}$$

Since $P_H(E^n)$ approaches 0, and $P_I(E^n)$ is bounded below by $P_I(G) > 0$, the odds $\frac{q_n}{1-q_n}$ must limit to zero as well. So $P(G \mid E^n)$ approaches $P_I(G \mid E^n)$, which by assumption approaches one. The second part follows since G and Q are incompatible.

We must distinguish between general *qualitative* evidence and evidence *de re*. The evidence that does the work in Proposition 1 is not simply that *there are* lots of things which are black-if-ravens. Each of the general qualitative propositions "At least n things are black-if-ravens" is entailed by Plenitude, and thus none of these propositions are evidence against Plenitude. Indeed, *no* qualitative proposition in the structure algebra which is consistent with Plenitude is evidence against Plenitude. This includes any *simple existential generalization* of the form $\exists x_1 \cdots \exists x_n \, \phi$ where ϕ is a quantifier-free formula without names. Such existential propositions will typically be evidence *for* Plenitude. (The likelihood of a simple existential given a Humean hypothesis is one, while its likelihood will be less than one for any "induction-friendly" hypothesis that allows the opposite generalization $\forall x_1 \cdots \forall x_n \, \neg\phi$ to be inductively confirmed.)

But qualitative propositions like these are not the only kind of evidence we have. The evidence that does the work in Proposition 1 consists of stronger *de re* propositions about particular things—*this* is black-if-a-raven, *that* is black-if-a-raven, and so on. (It's worth keeping in mind the basic point that even if E is evidence for Q, a proposition that *entails* E can be evidence *against* Q.) Furthermore, while Plenitude implies that there are many black ravens (and many non-black things), it does not imply that *this* is such a thing. And while Plenitude implies that many large regions are free of non-black ravens, it does not imply that *this* is such a region. So the evidence that this *is* such a region can disconfirm Plenitude. Even if you're sure there are lots of black ravens out there, it can still come as a surprise to discover them *here*. (See White 2000 for illuminating and closely related discussion.)

This might sound rather controversial. Many philosophers have engaged in a project of *reducing* attitudes *de re*—such as believing that *this* is a raven—to "self-locating" attitudes *de se*—such as self-ascribing the qualitative property of being appropriately related to something or other which is a raven (see for example, Lewis 1979; Chalmers 2011; Ninan 2013; for critical discussion see Cappelen and Dever 2013; Magidor 2015). This approach will extend to credences and evidence. One might report this in a slogan by saying that "all evidence is qualitative"—but one should not be misled by this to think that all evidence consists in *qualitative propositions* in the sense I have used the term. A qualitative proposition is purely about the general pattern of certain fundamental properties and relations (here modeled by a set of structures which is invariant under isomorphisms). A proposition like this is *de dicto* and "boring"—true once and for all, not the sort of thing that distinguishes some perspectives from others. But even if *de re* beliefs are really self-locating beliefs, they are surely not beliefs in such boring contents. Believing *this is a raven* may well be reducible to having some qualitative property, but it is not reducible to believing some general qualitative proposition of this kind. Two individuals may believe the very same general qualitative facts about the global structure of the world, while having different views about where they find themselves in that pattern, and in particular different views about whether they themselves are pointing at a raven. One of them believes *this is a raven,* while the other believes no such thing. Likewise, even if *de re* evidence is reducible to *de se* evidence, that wouldn't make it *qualitative* evidence in the sense of the term used here.

Furthermore, even if attitudes *de re* can be reduced to attitudes *de se,* this is not to say there is anything wrong with the framework I have used here. We can still model *de re* contents using sets of structures, and assign such contents probabilities. It's just that this model does not give the whole story. There is more to say about *how* the domain D represents possible individuals: ultimately, its elements shouldn't just stand for "bare" individuals, but rather for certain ways of being appropriately *related* to an individual.[25] Filling in these details of our theory of *de re* attitudes need not make any direct difference to what I have said about the import of *de re* evidence for Plenitude.

Need not—but still, the "self-locating turn" *might* make a difference indirectly, by motivating amendments to the standard account about what evidence supports. The last twenty years have seen a great deal of work on the question of how to update on self-locating information (for overview see Titelbaum 2016). Some philosophers think that what your self-locating evidence supports is not correctly captured by prior conditional probabilities in the standard Bayesian way I have taken for granted; according to the reductionist view we are discussing, this will extend to *de re* evidence such as instances of simple generalizations.

For example, Meacham (2008; see also Halpern 2004) proposes that we should update by "compartmentalized conditionalization." For any attitude-content A, let the *qualitative content* of A be the strongest qualitative proposition that is entailed by A. Let a *cell* be a maximal consistent qualitative proposition. The proposal is to update on evidence E by first conditionalizing on the qualitative content of E, and then conditionalizing on E within each cell. Here is an upshot of this proposal:

(Qual) For any qualitative hypothesis H, if E confirms H, then the qualitative content of E confirms H.

[25]We will want to say something like this anyway: we have been ignoring this complication, but because of the possibility of Frege puzzles, the elements of D should probably really be thought of as representing certain special *modes of presentation* for possible individuals, and not "bare" individuals themselves (see Footnote 14). We can think of the *de se* reductionist as giving a distinctive account of these modes of presentation.

Roughly speaking, this says that *merely* self-locating information cannot make a difference to qualitative hypotheses. So "compartmentalized conditionalization" conflicts with the moral I have drawn—namely, that while there is no *qualitative* proposition that disconfirms Plenitude, we can still have *non-qualitative* evidence that does disconfirm Plenitude.

On this point I take sides with some self-locating update schemes against others: (Qual) is false. One argument against it appeals to a principle of "considering the opposite":[26]

(Opp) If not-E disconfirms H, then E confirms H.

There is no "heads I win, tails we're even" evidence.[27] Here is a simple example to show how (Qual) conflicts with (Opp). Suppose you are sure there are exactly two people in the world, and your prior credence is evenly split between three hypotheses: *everyone is nice*, *only I am nice*, and *only the other person is nice*. It is clear that the evidence *I am not nice* disconfirms the qualitative generalization that everyone is nice. So by (Opp), *I am nice* should *confirm* the generalization that everyone is nice. But the qualitative content of *I am nice* is just *someone is nice*, which does not confirm the generalization. (It is entailed by all three hypotheses.)

Of course, there are still many different theories of self-locating evidence among those that reject (Qual), and the differences between them might still matter for the story I have told about Plenitude. Exploring this issue thoroughly would turn this article into a book. But I think that as long as we reject (Qual), the most important bit of the story is secure: *de re* evidence (whether conceived of as self-locating or not) can in principle disconfirm the qualitative hypothesis of Plenitude even if general qualitative evidence does not.

Still, *do* we in fact have the sort of evidence that disconfirms Plenitude? We have observed lots of instances of substantive regularities involving what may well be fundamental properties and relations. Have we observed enough of them to render Plenitude improbable, and justify incredulity? It's enough if we have very strong evidence for the truth of *any* logically contingent universal regularity—say, that every electron has mass, unrestrictedly speaking, or that there are no talking donkeys,

[26]For other arguments, see Bradley (2011, sec. 9) and Titelbaum (2016). Titelbaum (2013) calls a thesis very similar to (Qual) the *Relevance-Limiting Thesis*.

[27]This is not to commit to the claim that whenever it is possible to have not-E as evidence, it is also possible to have E as evidence. For example, the principle says that if *I am not dead* disconfirms *the gun was loaded*, then *I am dead* confirms *the gun was loaded*; but this is not to say that it is possible for me to ever have *I am dead* as part of my evidence.

unrestrictedly speaking. Of course we can't conclude that this is how our evidence stands just from technical results like Proposition 1. But at least we can see in principle how it might turn out this way. Putting things a bit roughly, the key precondition for disconfirming Plenitude—and justifying incredulous stares *a posteriori*—is the non-zero a priori probability of empirical inductive success.

But perhaps our inductive ambitions should be tempered. Here is a natural picture: some parts of reality might be too far away or isolated from us for us to observe. Then we should ask what conclusions we can draw for Plenitude if our evidence consists entirely of *observable* instances—intuitively, from around here. As far as Humean hypotheses go, this doesn't actually make much difference, as long as the observable domain is large. Given a Humean hypothesis H, even *restricted* simple generalizations are exceedingly unlikely to be true. First, suppose infinitely many instances are observable. Then the probability that they are all true given H is still zero (by the same argument). Given H, not only does unrestricted Plenitude have probability one, but in fact any infinite subdomain also has probability one of satisfying Plenitude. Second, if there is a large finite number of observable instances, the Humean probability that they are all true will still be very small, decreasing exponentially with the number of instances. So observing large finite regularities will still typically be strong evidence against H.

But while Plenitude is almost sure given a Humean hypothesis, the reverse may not be true. The best bet for Plenitude is if we reserve some prior credence for *non-Humean Plenitude.* Let's consider one more simple set-up. Suppose rational prior credence is split between three hypotheses: (1) a Humean hypothesis, **Humean Plenitude.** (2) An induction-friendly hypothesis, **Global Regularity,** conditional on which unrestricted simple generalizations have a non-zero chance of being true. (3) **Locally Regular Plenitude,** a *non-Humean* hypothesis according to which Plenitude is true—and thus all unrestricted simple generalizations are false—but also, conditional on which simple generalizations that are *restricted to observable instances* have a non-zero chance of being true. Locally Regular Plenitude says that reality is vast and various "out there," but may still be regular "around here." In this set-up, observing many instances of a simple generalization is *still* evidence against Plenitude—even though it is not evidence against Locally Regular Plenitude. Such observations confirm both Global Regularity and Locally Regular Plenitude, while disconfirming Humean Plenitude. As the number of observed instances increases, the probability of the Humean hypothesis will still approach zero—but it won't take Plenitude all the way there with it. Rather, in the limit the probability of Plenitude will diminish until it converges to that of Locally Regular Plenitude; this hypothesis was initially less probable than Plenitude, because it is more specific. So even in this case, where we have reserved some prior probability for the possibility that observable instances are not typical of the rest of reality, observing instances of regularities can still provide evidence against Plenitude (though this evidence is not as overwhelmingly strong as in the simpler set-up). We would only avoid this conclusion if one of Global Regularity or Humean Plenitude had rational prior probability *zero*.

Note also that the kind of empirical disconfirmation for Pattern Plenitude I have discussed does *not* straightforwardly extend to other kinds of "multiverse" theories. Observing many instances of a regularity does not disconfirm theories that say reality is vast and various, but still appropriately *regular* (and thus not utterly Plenitudinous). According to some multiverse theories from physics, there is a lot of stuff (Melbourne-counterparts and talking donkeys and so on) but it is all mundane physical stuff obeying logically contingent physical laws. Furthermore, if these laws comprise a *nice theory* T in the sense from Section 4, then the more restricted *T-Plenitude* would survive the kind of empirical disconfirmation just described, and *T-Humean hypotheses* may even gain empirical support along with the theory T.

Here's one last point worth noticing. At the outset, I took for granted a certain set L of "perfectly natural" or "fundamental" properties and relations. These have been treated with special privilege all along: we used them to specify the structure algebra, a "fundamental language," and the class of simple generalizations. Some philosophers have suggested that "metaphysical structure"—like Lewis's natural properties—might play a role in guiding rational inductive inference (Sider 2011, sec. 3.3). But any connection there might be between "metaphysical structure" and rational degrees of belief is not well-understood. This essay was occasioned by a question in the epistemology of metaphysics—how probable is Plenitude?—but I think that the philosophical exploration of *structure measures* is also a useful step in a project that we might call the "metaphysics of epistemology."

A Appendix

This appendix gives a precise statement and proof sketch for the result in Section 4: for any "nice" theory T, there is a "random T-model" that has probability one up to isomorphism, with respect to any T-Humean structure measure.

Definition 8. (Hodges 1997, 160) A first-order L-theory T is **nice** iff it has the following two properties:

(a) **The Hereditary Property.** Every substructure of a model of T is a model of T;

(b) **The Amalgamation Property**. For any three T-models, A, B and C, with embeddings $f : A \to B$ and $g : A \to C$, there are a T-model D and embeddings $f^+ : B \to D$ and $g^+ : C \to D$ such that $f^+ \circ f = g^+ \circ g$.

$$\begin{tikzcd} & D & \\ B \arrow[ur, dashed, "f^+"] & & C \arrow[ul, dashed, "g^+"'] \\ & A \arrow[ul, "f"] \arrow[ur, "g"'] & \end{tikzcd}$$

The first condition cuts structures down, and the second condition pastes structures together, possibly with overlap.

We can also characterize nice theories syntactically. An $\forall_1$ **formula** has the form $\forall x_1 \cdots \forall x_n\, \phi$ where ϕ is quantifier-free. (In Section 6 I called these *simple generalizations.*) An $\forall_1$ **theory** is a set of $\forall_1$ formulas. Similarly, an $\exists_1$ **formula** has the form $\exists x_1 \cdots \exists x_n\, \phi$ where ϕ is quantifier-free.

Theorem 2.

(a) T has the Hereditary Property iff T is logically equivalent to an $\forall_1$ theory. (By the Łoś-Tarski Theorem, see Hodges 1997, Theorem 5.4.4 on p. 143).

(b) T has the Amalgamation Property iff, for any $\forall_1$ formulas α_1 and α_2 such that $T \vdash \alpha_1 \vee \alpha_2$, there are $\exists_1$ formulas β_1 and β_2 such that

$$T \vdash \beta_1 \vee \beta_2$$
$$T \vdash \beta_1 \to \alpha_1$$
$$T \vdash \beta_2 \to \alpha_2$$

(Bryars 1973; see Bacsich and Hughes 1974, Corollary 2.5 on pp. 438–9)

Part (a) says that T has to consist entirely of simple universal generalizations. Part (b) says, roughly, that if T implies a disjunction of two different universal generalizations, then we can figure out that one disjunct is true just by observing some finite bit of the universe (a witness to one of the two existential generalizations β_1 or β_2).

Theorem 3. If T is a nice theory, then T has a countable model S such that

(a) S embeds every countable model of T;

(b) Any T-Humean structure measure (Definition 5) assigns the set of structures isomorphic to S probability one.

Theorem 1 in Section 3 is a special case of Theorem 3 (since the empty theory is nice).

Definition 9. We say S has the **Extension Property** for T iff for any finite T-models A and B such that A is a substructure of B, any embedding of A in S can be extended to an embedding of B in S.

$$\begin{tikzcd} A \arrow[r, "\subseteq"] \arrow[dr, "f"'] & B \arrow[d, dashed, "f^+"] \\ & S \end{tikzcd}$$

The Extension Property for a nice theory T implies that S embeds every countable T-model. (If S' is a T-model, each of its finite substructures is a T-model; it then follows that S embeds S' by Hodges 1997, Lemma 6.1.3, pp. 161–162.) Furthermore, any two countable structures with the Extension Property are isomorphic (Hodges 1997, Lemma 6.1.4, p. 162). So the following lemma suffices for Theorem 3.

Lemma 1. If T is a nice theory and P is a T-Humean structure measure, then the Extension Property for T has P-probability one.

Proof Sketch. The basic idea is that each finite extension of a finite structure has infinitely many opportunities to turn up, and the Humean condition guarantees that its probability of turning up never gets too small.

If A is a finite structure with domain $1, 2, \ldots, n$, and $a_1, \ldots, a_n$ are distinct elements of D, let $A(\bar{a})$ be the set of L-structures S with domain D such that the function $(n \mapsto a_n)$ is an embedding. Intuitively, this is a proposition that says that the individuals $a_1, \ldots, a_n$ instantiate the pattern A. Here is another way of saying that a structure S has the Extension Property for T: for any finite T-models A and B such that A is a substructure of B, for any individuals $a_1, \ldots, a_m$ such that S satisfies $A(\bar{a})$, there are distinct individuals $b_1, \ldots, b_n$ such that S also satisfies $B(\bar{a}, \bar{b})$.

Fix A, B, and $\bar{a}$, and enumerate an infinite sequence of tuples $\bar{b}^1, \bar{b}^2, \ldots$, such that all of the individuals in these tuples are distinct from one another. Let B_i abbreviate $B(\bar{a}, \bar{b}^i)$. Then (since there are only countably many such combinations of finite structures and tuples) it suffices to show that for each such A, B, and $\bar{a}$,

$$P(A(\bar{a})\ \neg B_1\ \neg B_2 \cdots) = 0$$

In other words, the probability is zero that the extended structure B does not show up anywhere. If there is some n such that $P(A(\bar{a})\ \neg B_1 \cdots \neg B_n) = 0$, then we're done. Assuming otherwise, it suffices to show that there is some $\delta > 0$ such that for each n,

$$P(B_{n+1} \mid A(\bar{a})\ \neg B_1 \cdots \neg B_n) > \delta$$

We can get this from the T-Humean condition, as follows.

Each of the propositions $A(\bar{a}), B_1, \ldots B_n$, which say that certain individuals instantiate a certain finite pattern, can be written as a finite conjunction of basic propositions. The Boolean combination $A(\bar{a})\ \neg B_1 \cdots \neg B_n$ can then be rewritten in disjunctive normal form, as a finite disjunction $D_1 \vee \cdots \vee D_k$ of mutually exclusive finite conjunctions of basic propositions. For each disjunct D_i that is consistent with T, the Amalgamation Property guarantees that B_{n+1} is also consistent with D_i and T. (If T has a model that satisfies D_i, the Amalgamation Property tells us that T also has an extended model that embeds B. This extended structure is still a model of $A(\bar{a})$ and D_i, and we can choose the extra elements of this structure to be those given by the tuple $\bar{b}^{n+1}$, so it is also a model of B_{n+1}.)

Write B_{n+1} as a finite conjunction of basic propositions $E_1 \cdots E_N$. The "complexity" N does not depend on n: it is determined by the number of elements of the structure B and the number of relations of each adicity in L. Each disjunct D_i with positive probability is consistent with T (since we have assumed $P(T) = 1$). In this case, E_j is consistent with $D_i E_1 \cdots E_{j-1}$ and T, so the T-Humean property tells us

$$P(E_j \mid D_i\, E_1 \cdots E_{j-1}) > \varepsilon$$

Multiplying the conditional probabilities together,

$$P(B_{n+1} \mid D_i) = P(E_1 \ \cdots\ E_N \mid D_i) > \varepsilon^N$$

Since this holds for each disjunct D_i with positive probability,

$$P(B_{n+1} \mid A(\bar{a})\ \neg B_1 \cdots \neg B_n) = P(B_{n+1} \mid D_1 \vee \cdots \vee D_k) > \varepsilon^N$$

This is what we needed to show. □

References

Armstrong, D. M. 1989. *A Combinatorial Theory of Possibility*. Cambridge University Press.

Bacsich, Paul D., and Dafydd Rowlands Hughes. 1974. "Syntactic Characterisations of Amalgamation, Convexity and Related Properties." *The Journal of Symbolic Logic* 39 (3): 433–51. https://doi.org/10.2307/2272886.

Bradley, D. J. 2011. "Self-Location Is No Problem for Conditionalization." *Synthese* 182 (3): 393–411. http://www.jstor.org/stable/41477639.

Bryars, David A. 1973. "On the Syntactic Characterization of Some Model Theoretic Relations." London.

Cameron, Peter J. 1997. "The Random Graph." In *The Mathematics of Paul Erdös II*, edited by Ronald L. Graham and Jaroslav Nešetřil, 333–51. Berlin, Heidelberg: Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-60406-5_32.

Cappelen, Herman, and Josh Dever. 2013. *The Inessential Indexical: On the Philosophical Insignificance of Perspective and the First Person*. Oxford University Press. https://www-oxfordscholarship-com/view/10.1093/acprof:oso/9780199686742.001.0001/acprof-9780199686742.

Carnap, Rudolf. 1950. *Logical Foundations of Probability*. Chicago]University of Chicago Press.

Cavendish, Margaret. (1666) 2001. *Observations Upon Experimental Philosophy*. Edited by Eileen O'Neill. Cambridge University Press.

Chalmers, David J. 2011. "Frege's Puzzle and the Objects of Credence." *Mind* 120 (479): 587–635.

de Moivre, Abraham. 1738. *The Doctrine of Chances: Or, a Method of Calculating the Probabilities of Events in Play*. The Second Edition: Fuller, Clearer, and More Correct Than the First. London: H. Woodfall. https://books.google.com/books?id=PII_AAAAcAAJ&printsec=frontcover&source=gbs_ge_summary_r&cad=0#v=onepage&q=independent&f=false.

deRosset, Louis. 2011. "On the Plurality of Worlds: David Lewis." *Humana Mente* 19.

Dorr, Cian. n.d. "Counterparts."

Erdős, Paul. 1959. "On Random Graphs I." *Publicationes Mathematicae (Debrecen)* 6: 290–97.

Erdős, P., and A. Rényi. 1963. "Asymmetric Graphs." *Acta Mathematica Academiae Scientiarum Hungarica* 14 (3): 295–315. https://doi.org/10.1007/BF01895716.

Forrest, Peter, and D. M. Armstrong. 1984. "An Argument Against David Lewis' Theory of Possible Worlds." *Australasian Journal of Philosophy* 62 (2): 164–68.

Halpern, Joseph. 2004. "Sleeping Beauty Reconsidered: Conditioning and Reflection in Asynchronous Systems." In *Proceedings of the Twentieth Conference on Uncertainty in Ai*, edited by Tamar Szabo Gendler and John Hawthorne, 1:111–42. Oxford University Press.

Hájek, Alan. 2003. "What Conditional Probability Could Not Be." *Synthese* 137 (3): 273–323.

Hodges, Wilfrid. 1997. *A Shorter Model Theory*. 1 edition. Cambridge ; New York: Cambridge University Press.

Howson, C. 1973. "Must the Logical Probability of Laws Be Zero?" *The British Journal for the Philosophy of Science* 24 (2): 153–63. https://doi.org/10.1093/bjps/24.2.153.

Hume, David. (1748) 2007. *An Enquiry Concerning Human Understanding and Other Writings*. Cambridge University Press.

———. (1739) 2007. *A Treatise of Human Nature: A Critical Edition*. Oxford University Press.

Kant, Immanuel. (1781) 1999. *Critique of Pure Reason*. Edited by Paul Guyer and Allen W. Wood. Cambridge; New York: Cambridge University Press.

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

Kraay, Klaas J. 2010. "Theism, Possible Worlds, and the Multiverse." *Philosophical Studies* 147 (3): 355–68.

Lewis, David. 1979. "Attitudes de Dicto and de Se." *Philosophical Review* 88 (4): 513–43.

Lewis, David K. 1986. *On the Plurality of Worlds*. Wiley-Blackwell.

———. 2009. "Ramseyan Humility." In *Conceptual Analysis and Philosophical Naturalism*, edited by David Braddon-Mitchell and Robert Nola, 203–22. MIT Press.

Libkin, Leonid. 2013. *Elements of Finite Model Theory*. Springer Science & Business Media.

Lovejoy, Arthur O. 1936. *The Great Chain of Being: A Study of the History of an Idea*. Transaction Publishers.

Magidor, Ofra. 2015. "The Myth of the de Se." *Philosophical Perspectives* 29 (1): 249–83.

McHarry, John D. 1978. "A Theodicy." *Analysis* 38 (3): 132–34.

Meacham, Christopher J. G. 2016. "Ur-Priors, Conditionalization, and Ur-Prior Conditionalization." *Ergo: An Open Access Journal of Philosophy* 3.

———. 2008. "Sleeping Beauty and the Dynamics of de Se Beliefs." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 138 (2): 245–69. http://www.jstor.org/stable/40208872.

Ninan, Dilip. 2013. "Self-Location and Other-Location." *Philosophy and Phenomenological Research* 87 (1): 301–31.

Nolan, Daniel. 1996. "Recombination Unbound." *Philosophical Studies* 84 (2-3): 239–62.

Noonan, Harold W. 2014. "The Adequacy of Genuine Modal Realism." *Mind* 123 (491): 851–60.

Peterman, Alison. 2017. "Empress Vs. Spider-Man: Margaret Cavendish on Pure and Applied Mathematics." *Synthese*, July. https://doi.org/10.1007/s11229-017-1504-y.

Popper, Karl R. (1934) 1959. *The Logic of Scientific Discovery*. Routledge.

Russell, Jeffrey Sanford, and John Hawthorne. 2018. "Possible Patterns." *Oxford Studies in Metaphysics* 11.

Saucedo, Raul. 2011. "Parthood and Location." In *Oxford Studies in Metaphysics: Volume 5*, edited by Dean Zimmerman and Karen Bennett. Oxford University Press.

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

Tegmark, Max. 2008. "The Mathematical Universe." *Foundations of Physics* 38 (2): 101–50.

Titelbaum, Michael G. 2016. "Self-Locating Credences." In *The Oxford Handbook of Probability and Philosophy*, edited by Alan Hajek Christopher Hitchcock. Oxford University Press.

———. 2013. "Ten Reasons to Care About the Sleeping Beauty Problem." *Philosophy Compass* 8 (11): 1003–17. https://doi.org/10.1111/phc3.12080.

Wallace, David. 2012. *The Emergent Multiverse: Quantum Theory According to the Everett Interpretation*. Oxford University Press.

Wang, Jennifer. 2013. "From Combinatorialism to Primitivism." *Australasian Journal of Philosophy* 91 (3): 535–54.

White, Roger. 2000. "Fine-Tuning and Multiple Universes." *Noûs* 34 (2): 260–76.

———. 2018. "Reasoning with Plenitude." In *Knowledge, Belief, and God: New Insights in Religious Epistemology*, edited by Matthew A. Benton, John Hawthorne, and Dani Rabinowitz. Oxford: Oxford University Press.

Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford University Press.

———. 2013. *Modal Logic as Metaphysics*. Oxford University Press.

Wilson, Jessica M. 2010. "What Is Hume's Dictum, and Why Believe It?" *Philosophy and Phenomenological Research* 80 (3): 595–637.

Yagisawa, Takashi. 1992. "Possible Worlds as Shifting Domains." *Erkenntnis* 36 (1): 83–101.