

ISLAND UNIVERSES AND THE ANALYSIS OF MODALITY

Phillip Bricker

University of Massachusetts, Amherst

Is physical reality a single, unified whole – say, a spacetime of four, or ten, or twenty-six dimensions? Or does physical reality divide into parts that are causally and spatiotemporally isolated from one another, so-called “island universes”?¹ I doubt we could ever have good reason to answer this question one way or the other. In this paper, I focus instead on a prior question, one more amenable to philosophical analysis: Is the supposition that island universes exist even coherent? Is it metaphysically or logically possible? I will argue that island universes are metaphysically possible, and then consider the substantial impact this has upon realist theories of possible worlds and the realist analysis of modality.

The question whether island universes are possible has been raised intermittently through the history of philosophy, and most often answered in the negative. Kant maintained the necessary unity of space and of time: necessarily, everything spatial is spatially related to everything else spatial; *mutatis mutandis* for time.² Bradley disagreed about the unity of space and of time, but held that reality (the Absolute) is necessarily unified in some non-spatial, non-temporal way.³ In this century, logical positivists held that the supposition that island universes exist is meaningless – and so, *a fortiori*, not genuinely possible – on grounds of the verifiability criterion of meaning: the existence of a part of physical reality causally and spatiotemporally isolated from us would be unverifiable (by us), even in principle.⁴ I have little sympathy with verificationist or idealist arguments against the possibility of island universes, but nothing new to say against them. I mention them to set them aside.

More recently, the possibility of island universes has been challenged by David Lewis.⁵ His argument is indirect. Island universes make trouble for his realist theory of possible worlds; to the

¹In astronomical usage, galaxies have been called “island universes” owing to their relative isolation. In metaphysical usage, the isolation of “island universes” is more complete. The phrase ‘island universe’, in its metaphysical sense, occurs in Armstrong (1989a). (See § 2.1 below for a precise definition.)

² Kant (1781), A25, B39. For critical discussion, see Strawson (1966).

³Bradley (1893). For arguments against the unity of space and of time, see pp. 210-14, 286-8; for arguments that reality (the Absolute) is necessarily unified, see pp. 140-3.

⁴ In arguing against logical positivism, Russell ponders whether the supposition that “there is a cosmos which has no spatiotemporal relation to the one in which we live” is meaningful, even though there could be no evidence for or against it. See Russell (1950), p. 278.

⁵ See Lewis (1986), pp. 69-78.

extent that he has good and sufficient reason to accept the theory, he has good and sufficient reason to reject the possibility of island universes. Trouble comes when one tries to say, as any realist must, how the possible worlds are demarcated one from another. According to Lewis, possible individuals are part of one and the same possible world if, and only if, they are spatiotemporally related.⁶ It follows immediately that no possible world is composed of island universes, of spatiotemporally isolated parts. Given the standard analysis of possibility as truth at some possible world, island universes, then, are impossible.

I agree that island universes make trouble for Lewis's realism about possible worlds, perhaps more than Lewis realizes. But Lewis's realism is not the only realism in town. On Lewis's brand of realism, the actual world and the merely possible worlds are ontologically on a par: there is no absolute ontological distinction between the actual and the merely possible. If, instead, the realist endorses some form of absolute actuality, then, I will argue, there is a simple and natural solution to the problem of island universes. Simple and natural, but in one sense radical, because it involves an emendation in the standard analysis of possibility as truth at some possible world, and in the standard method of giving truth conditions relative to possible worlds. On the emendation, as I will develop it, modal operators are analyzed as *plural*, rather than *individual*, quantifiers over possible worlds.

When the amended analysis of possibility is combined with absolute actuality, a number of problems faced by Lewis's realism are neatly solved. (1) The possibility of island universes, which I defend in §2.4 and §2.6, can easily be accommodated. (2) The possibility of nothing, which I defend in §4.2, can also easily be accommodated (if so desired). (3) A version of Lewis's "principle of recombination" – roughly, that anything can coexist with anything – can be accepted without qualification; it is invulnerable to the Forrest-Armstrong argument (§4.1). (4) The principle of the identity of qualitatively indiscernible worlds, mysteriously undecidable on Lewis's theory, can be decisively refuted (§4.3).⁷ All this is good news for those of us who are favorably inclined towards realism about possible worlds, and the project of analyzing modality in terms of them. Those who reject realism about possible worlds on grounds of crazy or bloated ontology, and who have no misgivings about primitive modality, will not, of course, be moved.

1. REALIST THEORIES

1.1. *Realism*. What is realism about possible worlds? It won't do simply to say that, according to the realist, possible worlds exist: given the range of disagreement over what sort of thing 'possible world' refers to, those who assert that possible worlds exist cannot be taken to share a

⁶ Actually, Lewis holds that possible individuals are part of the same world if and only if they are spatiotemporally related, or *analogically* spatiotemporally related. (See Lewis (1986), pp. 74-76.) Since this complication won't matter for what follows, I will simply use 'spatiotemporal' broadly so as to include what Lewis calls "analogically spatiotemporal."

⁷ All four difficulties with Lewis's view are discussed in Armstrong (1989a); the second difficulty, according to Armstrong, is not genuine.

single view. Nor will it do to say that, according to the realist, possible worlds exist and are *concrete*: there is too much disagreement (and unclarity) over what to call “concrete.”⁸ The following characterization, however, should be sufficiently definite for our purposes. Call a theory of possible worlds *realist* if it holds that (1) worlds exist⁹; (2) worlds are individuals rather than classes, or functions, or mathematical models; (3) worlds are particulars rather than properties or universals; (4) (most) worlds are complex rather than simple – for example, many worlds have parts that stand in spatiotemporal relations to one another. The problems I raise in this paper involving island universes are problems for realism, so characterized. Thus, I address my argument and proposed solution primarily to realists. But non-realists who have as their goal an ontologically “innocent” reinterpretation of realist theory – various fictionalists and ersatzers – will also want to take heed.

1.2. *Actuality*. According to David Lewis’s brand of realism about possible worlds, actuality is “indexical” and relational: the actual world is *this* world, the world *we* inhabit; to be actual is to be appropriately related to *us*. And that is all. *Being actual* confers no special ontological status; our world and the other possible worlds do not differ in ontological kind. The difference, so to speak, is that we are here and not there.¹⁰

The alternative for the realist is to hold that actuality is absolute, that there is an ontological distinction of kind between the actual and the merely possible. In my opinion, this is the only viable option for the realist. Our conceptual scheme demands that actuality be *categorical*: whatever is of the same ontological kind as something actual is itself actual. To hold, then, as Lewis does, that the actual world and the merely possible worlds do not differ in kind is simply incoherent.¹¹ Our conceptual scheme also demands, I believe, a partial alignment between ontological and epistemological distinctions. In particular, whatever is of the same ontological kind as (and external to) a thinking subject cannot be known *a priori* to exist. To hold, then, as Lewis does, *both* that the actual world and the other possible worlds do not differ in kind *and* that we know *a priori* that other possible worlds exist is, again, simply incoherent.¹² These are not

⁸ See Lewis (1986), pp. 81-5.

⁹ I use ‘world’ and ‘possible world’ interchangeably. I use ‘exist’ here to range over absolutely everything, the broadest category of being. In what follows, it will be clear from context, I hope, whether or not ‘exist’ is to be restricted to actual things.

¹⁰ See Lewis (1973), pp. 84-91; Lewis (1986), pp. 1-3, 92-6.

¹¹ One must be careful not to conflate two related objections: the objection that actuality must be *categorical*; and the objection that actuality must be *universal*, applying to whatever has being, in the broadest sense of ‘being’. Lewis (1986) discusses only the latter – I think, less plausible – objection. Note that either objection could be behind the tendency of philosophers to interpret Lewis – his denials notwithstanding – as believing in an actuality “much bigger and more fragmented than we ordinarily think.” See Lewis (1986), pp. 97-101.

¹² This objection was raised early on in Skyrms (1976), p. 326.

arguments, of course. But, if I am right, no arguments are needed: the alleged incoherence is right on the surface.

A coherent realism, then, must hold that the actual and the merely possible differ in kind. I would develop the view, within a broader context, as follows. There are two fundamental categories of being. There is *physical* reality, which contains us as a part, and which we have epistemic access to, when we do, through causal interaction. And there is *metaphysical* reality, which contains the possible worlds (and, I think, mathematical entities and structures) as parts, and which we have epistemic access to, when we do, through our powers of thought, through our mental representations. The distinction I have in mind between physical and metaphysical does not coincide, on current usage, either with the distinction between actual and non-actual, or with the distinction between concrete and abstract. The mathematical part of what I call metaphysical reality – if not the whole of it – is typically called “actual” (by realists and non-realists alike); it would be mistaken, therefore, to identify physical reality with “actuality.” The modal part of metaphysical reality, as I conceive it, may properly be called “concrete,” since it is qualitatively no less determinate than physical reality; it would be misleading, therefore, to identify physical reality with the “concrete.”¹³

On the brand of realism I am propounding, there is an absolute fact as to which among all the possible worlds has been actualized; call it realism *with absolute actualization*. Absolute actualization comes in two versions depending upon whether the whole of physical reality is itself one of the possible worlds, or instead has a representative among the possible worlds. To make the different versions vivid, consider Leibniz’s God surveying the possible worlds prior to actualization.¹⁴ Did actualization consist in God conferring a special ontological status upon the best possible world? Strictly speaking, this is transformation, not creation. Or did actualization consist in God creating *ab nihilo* a new “world,” qualitatively indistinguishable from the best possible world, but differing from it in ontological status? I think the second version is superior, but I won’t argue that here; the difference won’t much matter for what follows.¹⁵

¹³ My distinction between physical and metaphysical is closer to the outmoded distinction between existence and subsistence. Note that I do not exclude irreducibly mental entities – if any there be – from physical reality.

¹⁴ Here and throughout, I call upon Leibniz’s God to illustrate absolute actualization. But there are two implications of the image that I want to explicitly reject: I think actualization is primitive as well as absolute; and I deny that God could actualize a world without being a part of the world He actualized.

¹⁵ On this second version, what we call “the actual world” – that is, physical reality – is not, technically speaking, a world. But use of ‘the actual world’ is so ingrained that I won’t go out of my way to avoid it. The second version of realism with absolute actualization falls under what Lewis calls “pictorial ersatzism” (Lewis (1986), pp. 165-174). But its realist credentials are not thereby impugned. Indeed, Lewis’s chief criticism of pictorial ersatzism is that it offers no ontological gain for those who find his own account of worlds incredible. Agreed. The gain lies

No doubt I will be asked to explain further this distinction between physical and metaphysical reality. I can say something the distinction is *not*: it is not a difference in some quality. For any quality, there would have to be worlds at which some things have, and some things lack, the quality. But at every possible world everything is metaphysical, and at the actual “world” everything is physical. I can say something about how the distinction relates to *us*. What we think about, in one sense, is not limited by what actually exists – as when we think about golden mountains, or perfect cubes. Metaphysical reality is what we can think about, in that sense; it is the realm of “content” for all actual and possible thought. Physical reality, crudely put, is what we can bump into. I cannot, however, say anything about how physical and metaphysical reality differ in and of themselves. That is only to be expected: the distinction is fundamental and unanalyzable.

The most serious objection to realism with absolute actualization, as Lewis has emphasized, is that it seems to allow for a coherent skepticism about one’s own actuality, whereas such skepticism is absurd.¹⁶ Granted, such skepticism is absurd. When asked – how do I know that I am actual? – I can give only one response: I just know it. I think that response is acceptable when dealing with a fundamental ontological category; talk of “evidence” here is beside the point. Moreover, anyone who accepts more than one fundamental ontological category, be it individuals and classes, or particulars and universals, must face the same sort of question, and, I claim, give the same answer. How do I know that I am an individual and not a class? I just know it. How do I know that I am a particular and not a universal? I just do. If this is right, then all ontological pluralists are in the same boat. And ontological monism, while attractively simple, faces formidable obstacles.

I do not say that realism about possible worlds with absolute actualization is without difficulties. I wouldn’t stake my life on its being true – or even my next paycheck. But I can say with confidence that it is a serious contender in the metaphysical arena, worth seriously pursuing.

1.3. *Isolation and Unification: The Realist Analysis of World.*¹⁷ Realists, I have noted, must provide a criterion of demarcation for worlds. Moreover, this must be done without relying on modal notions lest the analysis of modality in terms of worlds run in a circle. The basic strategy is this. Let *logical space* be that part of metaphysical reality over which (alethic) modal operators range; in other words, the aggregate sum of *possibilia*.¹⁸ Some regions of logical space are

elsewhere: in giving an adequate account of actuality, and in solving the problems addressed in this paper.

¹⁶ Lewis (1970); Lewis (1986), p. 93.

¹⁷ This section is adapted from Bricker (1996).

¹⁸ If logical space is a *proper* part of metaphysical reality, as I think, it is also imperative that the extent of logical space be characterized without relying on modal notions. I believe that can be done, but I won’t discuss it here.

unified; the maximal such unified regions are the *worlds*. On this basic strategy, I am in agreement with Lewis. But I disagree with Lewis over two substantial issues having to do with the manner of unification. First, for Lewis, all worlds are *globally* unified (or *almost globally* unified): at any world, (almost) every part is directly linked to (almost) every other part.¹⁹ I hold instead that some worlds are *locally* unified: at some worlds, parts are directly linked only to “neighboring” parts. Second, for Lewis, each world is *spatiotemporally* unified; every world is *spatiotemporally* isolated from every other. I hold instead: a world may be unified by non-spatiotemporal relations; every world is *absolutely* isolated from every other. If I am right with respect to (either or both) of these issues, then Lewis’s conception of logical space is impoverished: some possible worlds are missing. I have argued for my views elsewhere; here I present the requisite definitions, and my analysis of world, without defense.²⁰

First, I need to introduce some familiar metaphysical machinery.²¹ I start with an abundant conception of property according to which, for every class of *possibilia*, there is (at least) one property had by all and only the members of that class. Then I need the distinction between the *fundamental*, or (perfectly) *natural* properties and relations, and the rest. The natural properties and relations are those that correspond to immanent universals or tropes, if there are universals or tropes. They make for qualitative similarity: if two things instantiate the same natural property, or each divides into parts that stand in the same natural relation, then the things are objectively similar in some qualitative respect. Moreover, the qualitative supervenes on the natural: fixing the natural properties and relations suffices to fix all the qualitative properties and relations.

In terms of naturalness, a number of indispensable metaphysical notions can be defined. I will be brief.²² Things are (intrinsic, qualitative) *duplicates* just in case there is a *similarity map* from one to the other: a one-to-one correspondence between their parts that preserves all natural properties and relations (and the part-whole relation). An *intrinsic nature* is a property had by all and only the duplicates of some thing. An *intrinsic* property is one that never differs between duplicates; a property is *extrinsic* just in case it is not intrinsic. An *internal* relation is a relation

¹⁹ This follows from Lewis’s claim that the unifying relations are “pervasive.” See Lewis (1986), p. 76.

²⁰ In Bricker (1993), I argue that if Einsteinian relativity is true (on its most natural interpretation), then we live in a *locally* unified world; such worlds had better be possible! In Bricker (1996), I defend the view that worlds are *absolutely* isolated from one another. (Note, however, that for a realist who accepts universals – unlike myself – that view will have to be qualified.)

²¹ In this and the following paragraph, I more or less follow Lewis (1986), pp. 59-63.

²² Quantifiers range over all parts of physical and metaphysical reality unless explicitly restricted. I assume familiarity with mereology, the theory of part and whole. In particular, I assume unrestricted mereological composition: for any things, there is a (mereological) *sum*, or *fusion*, of those things.

that supervenes on the intrinsic natures of its *relata*. Having-the-same-mass-as is an example of an internal relation, assuming the mass properties are intrinsic. An *external* relation is one that, although it fails to supervene on the intrinsic natures of its *relata*, does supervene on the intrinsic natures of its *relata*, and of the fusion of its *relata*.²³ Being-adjacent-to is an example of an external relation: whether two things are adjacent to one another is not determined by their intrinsic natures, taken separately, but it is determined if one also takes into account the intrinsic nature of their fusion. A relation that is either internal or external is *intrinsic*; all others are *extrinsic*. Note that it is built into the definitions that all natural properties and relations are intrinsic.

I am ready, finally, to define unification, and then in terms of unification to analyze the notion of world. Unification, I take it, is to be characterized in terms of the holding or failing to hold of natural external relations. But which relations? Different choices yield different notions of unification. I will give the definitions simultaneously with respect to two choices: the natural spatiotemporal relations, and all natural external relations. Two parts of logical space are *spatiotemporally separated* (*externally separated*) if and only if they are non-overlapping and no part of one stands in any natural spatiotemporal (external) relation to any part of the other. A part of logical space is *spatiotemporally unified* (*externally unified*) if and only if it is not the sum of two spatiotemporally separated (externally separated) parts. Two parts of logical space are *spatiotemporally related* (*externally related*) if and only if some spatiotemporally unified (externally unified) part of logical space includes them both; otherwise they are *spatiotemporally isolated* (*externally, or absolutely, isolated*).²⁴

Now we can analyze world in terms of unification. For Lewis, were he to accept the above definitions, the analysis would be this: a *world* is any maximal spatiotemporally unified part of logical space, that is, a spatiotemporally unified part not properly included in any other spatiotemporally unified part. As already noted, I reject this analysis as too narrow. (Indeed, for all we know, not even the actual “world” is spatiotemporally unified; perhaps, as physicists have pondered, spatiotemporal relations do not apply at the “sub-microscopic” level.) I accept instead: a *world* is any maximal externally unified part of logical space. From this it follows that worlds are absolutely isolated from one another, and, in particular, that no two worlds overlap. Finally, parts of logical space are *worldmates* if and only if they are part of the same world, if and only if

²³ More precisely, say that an ordered pair $\langle a, b \rangle$ and an ordered pair $\langle c, d \rangle$ are *internal duplicates* iff a is a duplicate of c and b is a duplicate of d ; *external duplicates* iff, in addition, the composite of any similarity map from a to c and any similarity map from b to d induces a similarity map from the fusion of a and b to the fusion of c and d . Then, an *internal* (dyadic) relation is one, the holding of which never differs between pairs that are internal duplicates; an *external* (dyadic) relation is one that is not internal, but the holding of which never differs between pairs that are external duplicates. (Analogously for relations of three or more places.)

²⁴ In Bricker (1996), I (unwisely) used ‘isolation’ both for what I here call ‘separation’ and (in informal discussion) for what I here call ‘isolation’. When worlds are not globally unified (see below), the difference matters.

they are externally related. In what follows, I will assume that worlds are demarcated by external interrelatedness, although all of my main arguments could be modified to apply to Lewis's notion of world.

I conclude this section with an example that serves to illustrate the definitions, and that will play a role in the argument of §2.4. Worlds may be unified to a greater or lesser degree. At one end of the spectrum, we have *globally unified* worlds at which no part is externally separated from any other part. At a globally unified world, points of spacetime (if such exist and are atomic) are directly linked to one another by some natural external relation, presumably, by some external relation of spatiotemporal distance (interval). At the other end of the spectrum, we have *locally unified* worlds at which the only parts that are *not* externally separated are overlapping or adjacent parts.²⁵ (The separated parts are nonetheless externally *related* in virtue of belonging to a single externally unified region of logical space.) At a locally unified world with continuous spacetime, distinct points of spacetime are separated (being non-adjacent), and so are not directly linked by any natural external relation; relations of spatiotemporal distance are *extrinsic*, rather than external, because the distance between points depends upon the intervening spacetime, upon the lengths of paths from one point to the other.²⁶

2. ARE ISLAND UNIVERSES POSSIBLE?

2.1. *Island Universes: Strong and Weak Senses.* We now have the resources to define the notion of island universe, and to consider whether or not island universes are metaphysically possible. But we must be careful to distinguish an *absolute*, or *strong*, sense of 'island universe' from various *non-absolute*, or *weak*, senses, depending upon whether island universes are required to be absolutely isolated, or only isolated in one or another respect. Thus, island universes in the *strong* sense exist if and only if physical reality is not externally unified; the island universes are the maximal externally unified parts of physical reality. Island universes in a (prominent) *weak* sense exist if and only if physical reality is not spatiotemporally unified; the island universes are the maximal spatiotemporally unified parts. Other weak senses of 'island universe' can be had by replacing 'spatiotemporally' with 'spatially' or 'temporally' or 'causally' or any combination of these.²⁷

²⁵ Topologically speaking, two regions are *adjacent* iff they are non-overlapping, but one contains a boundary point of the other. (For example, on the real line, the open interval (0, 1) is adjacent to the closed interval [1, 2], but not to the open interval (1, 2).) Only worlds with topological structure can be locally unified.

²⁶ In Bricker (1993), I argue that distance relations are extrinsic, rather than external, at (some) worlds with continuous spacetime.

²⁷ For Lewis, adding causal isolation to spatiotemporal isolation would have been redundant. See Lewis (1986), pp. 78-81. Not so for those who take causation to be a primitive external relation between events.

If worlds are demarcated by external interrelatedness, as I claim, then the possibility of island universes in a weak sense is not problematic; for example, a world might well have spatiotemporally isolated parts that are externally related by some non-spatiotemporal relation. The possibility of island universes in the strong sense, however, is blatantly contradictory (if possibility is truth at some world): worlds, being externally unified, cannot have absolutely isolated parts. In §2.4 and §2.6 I give my arguments for the possibility of island universes in the strong sense; in §3.1 through §3.5, I consider how the contradiction is to be avoided.

2.2. Lewis on the Possibility of Island Universes. If worlds are demarcated by spatiotemporal interrelatedness, as Lewis holds, then even the possibility of spatiotemporally isolated island universes is contradictory (if possibility is truth at some world): no world can both be spatiotemporally unified and have spatiotemporally isolated parts. Lewis therefore rejects the possibility of spatiotemporally isolated island universes (what he calls, “disconnected spacetimes”). He writes:

I would rather not [reject the possibility]; I admit some inclination to agree with it. But it seems to me that it is no central part of our modal thinking, and not a consequence of any interesting general principle about what is possible. So it is negotiable.²⁸

To make the rejection of island universes more palatable, Lewis offers up an assortment of substitutes. Perhaps, he suggests, when we think we conceive of the possibility of island universes, what we really have in mind is a big world, spatiotemporally unified, but with many causally isolated world-like parts. For example:

The spacetime of the big world might have an extra dimension. The world-like parts might then be spread out along this extra dimension, like a stack of flatlands in three-space.²⁹

And he provides three other ways that a big world might have causally isolated world-like parts.

I have no objection to Lewis’s substitutes: they are genuinely possible, one and all. But they do nothing to counter our inclination to believe in the possibility of island universes. If the notion of island universe were obscure, or very complex, it might be otherwise; we might not know what possibility we had in mind. But since the notion of island universe, once disambiguated, is simple and clear, Lewis’s substitutes are plainly beside the point. (Compare the question whether there could be a world at which space is “curved.” To point out the possibility of a world at which a “curved” space is embedded within a higher-dimensional “flat” space is unresponsive; it merely changes the subject.)

²⁸ Lewis (1986), p. 71.

²⁹ Lewis (1986), p. 72.

An “inclination to believe,” however, by itself, does not count for much. Lewis’s rejection of the possibility of spatiotemporally isolated island universes would be unassailable if, as he says, “[it is] not a consequence of any interesting general principle about what is possible.” I will argue below, on the contrary, that the possibility of island universes does follow from general principles about what is possible, from (strong versions) of the Humean denial of necessary connections. These principles, I think, will be hard for an arch-Humean such as Lewis to resist!

2.3. *Are Island Universes Physically Possible?* Before turning to metaphysical arguments for the possibility of island universes in the strong sense, I want to briefly consider some arguments loosely based upon contemporary physics. I grant, of course, that if island universes are physically possible, then, *a fortiori*, they are metaphysically possible. But arguments for the physical possibility of island universes either fail outright, or, in the best case, rest upon a controversial interpretation of objective chance. Thus, I rate the case for physical possibility, at best, as inconclusive.

To begin, consider the following proposals, all made by reputable physicists. (1) Perhaps our “universe” is but one of many “universes” in a series of “cosmic oscillations”: big bang, followed by big crunch, followed by big bang, followed by big crunch, and so on. The different “universes” may even differ in their physical laws. But, surely, if talk of “oscillations” is to be appropriate, there must be (spatio)temporal (and causal) relations between the many “universes”. (2) Perhaps our “universe” is the result of a quantum vacuum fluctuation, one of those things, according to quantum electrodynamics, “that happens from time to time.” And then why not more than once, resulting in many “universes”? But, surely, talk of “vacuum fluctuations” requires a pre-existing “vacuum,” and the many “universes” (for all the theory says) are spatiotemporally (and causally) related to one another via this “vacuum.” (3) Perhaps on the correct interpretation of quantum mechanics our “universe” is but one of many branching “universes” splitting off from one another with every measurement-like interaction. But since any two branching “universes” have a common history, events from different “universes” are spatiotemporally (and causally) related to one another via events that occurred before the “universes” split off. None of these proposals, then, provides an argument for the possibility even of spatiotemporally (and causally) isolated island universes, let alone island universes in the strong sense.³⁰

Perhaps General Relativity can do better. It won’t do, however, to simply note that Einstein’s field equation admits solutions with disconnected spacetimes because sums of solutions are solutions. Physicists do not automatically suppose that mathematical solutions to their equations have any genuine physical significance; further argument is needed. Try this. There is a well-known solution to Einstein’s field equation in which two regions of spacetime are connected

³⁰ See Leslie (1989), ch. 4, for a discussion (with references) of these and other physical mechanisms for generating many “universes.” If the many universes are needed only to solve the “fine tuning problem” – the problem of rendering unsurprising the fact that our “universe” is “fine tuned” for the existence of life – the “universes” needn’t be spatiotemporally (or causally) isolated; there need only be enough variety.

only by a single, short-lived “wormhole”: the “wormhole” evolves spontaneously between two spatially disconnected regions, it gradually grows to a maximum diameter, then it shrinks until, pinching off, it leaves behind two spatially disconnected regions as before. Thus far, we have only spatial, not spatiotemporal, isolation.³¹ But now consider physical indeterminism. Presumably, there was some non-zero objective chance that the “wormhole” would never have evolved, that is, some non-zero objective chance that the two regions would have been spatiotemporally isolated island universes. And, whatever has a non-zero objective chance of happening is physically possible.³²

I don't really know whether our best physical theory allows general relativity and indeterminism to be combined in the way required for this argument. In any case, I suspect the argument supports only the possibility of spatiotemporally isolated island universes, not the possibility of island universes in the strong sense. Here's why. In the possibility being envisaged, the spatiotemporally isolated island universes are *related* in virtue of their being some objective chance that a “wormhole” connect *one* with *the other*. This cannot be an *internal* relation, lest there also be some objective chance that a “wormhole” connect the one island with each and every duplicate of the other island, spread out through logical space! It must then be an *external* relation, either itself natural, or grounded in other natural external relations between the islands. The two islands, then, are externally related; they are not island universes in the strong sense.

2.4. *Island Universes and Plenitude: The Principle of Solitude (PS)*. So much for arguments based upon speculative physics. I turn now to arguments based upon speculative metaphysics. As already noted, I think the possibility of island universes (in the strong sense) follows from general principles of the plenitude of possibilities, in particular, from strong versions of the Humean denial of necessary connections. To start, let us consider what I call the *principle of solitude*, roughly: anything can exist all by itself. This principle captures the idea that actualization is unconditional: whether or not one thing is actualized is not conditional upon whether or not any other thing is actualized. Unconditional actualization, it seems to me, is part

³¹ Weingard (1976) uses this solution – due to Martin Kruskal – to argue that “it is physically possible for space (relative to a frame of reference) to be, at some time, in disconnected pieces.” (p. 220).

³² The idea of using objective chance at this point in the argument I get from Bigelow and Pargetter (1987). A related argument goes like this. Presumably, whether and how the “wormhole” evolved might depend upon features of the two spacetime regions, so that, had the regions been appropriately different, no “wormhole” would have evolved; the regions would have been spatiotemporally isolated. But, here, an appeal must be made to some principle of recombination to ensure that the regions *could* have been appropriately different; one may no longer be within the realm of *physical* possibility. In Bigelow and Pargetter (1987), the counterfactual argument is not distinguished from the argument from objective chance. Incidentally, in Lewis (1992), Lewis gives the counterfactual argument, and claims that “the intuitive case that island universes are possible has been much strengthened” thereby. But he does not suggest any realist response.

and parcel of the conception of intrinsic nature presented above (in §1.3): each thing has an intrinsic nature in virtue of which it can be conceived of apart from anything else; and, if it can be conceived apart from anything else, then, possibly, it exists apart from anything else, that is, all by itself.

To illustrate the principle of solitude, consider a discriminating God. While surveying the possible worlds prior to actualization, He comes upon the world corresponding to our actual “world.” He is not pleased with everything He sees. Perhaps only one thing pleases Him – say, Leibniz. Then, according to the principle of solitude, God could choose to actualize Leibniz, and nothing else. Or perhaps only one thing displeases Him – say, Hume. Then, according to the principle of solitude, God could choose to actualize the world minus only Hume, leaving a hole where Hume would have been. In deciding what to actualize, God does not have to take a world all or nothing: He can exercise a line item veto.

How should the principle of solitude be expressed within a realist framework? First, there is no reason to restrict the quantifiers to actual individuals; we can quantify universally over individuals from all the worlds. It is, however, restricted to worldbound individuals, individuals that are part of some world.³³ Second, since the principle is not a claim about the essences, or potentialities, of things, it is to be interpreted in terms of duplicates, rather than counterparts. The principle requires that a *duplicate* of Leibniz could exist all by itself; it does not say that the duplicate represents anything *de re* of Leibniz, or would properly be called “Leibniz.” (Perhaps some of Leibniz’s essential properties are extrinsic, having to do with his origins, or what have you.) Thus, the principle can be formulated as follows:

(PS) *Principle of Solitude*. For any worldbound individual, possibly, a duplicate of that individual exists all by itself.

Three brief comments may help to forestall misunderstanding. First, to say that the duplicate “exists all by itself,” of course, is to say that nothing wholly distinct from the duplicate exists: nothing exists but the duplicate *and its parts*. Second, (PS) does not rule out the necessary existence of “mathematical” entities: necessary beings are compatible with (PS) as long as they are not parts of worlds, and so not within the range of the quantifiers. Third, existing all by oneself is not the same as existing surrounded by empty spacetime. In the possibility posited by (PS), spacetime, if it exists at all, has whatever shape the duplicate has; in the case of lonely Leibniz, spacetime has the shape of a spacetime worm.

Now we are ready for the argument that (PS) leads to the possibility of island universes (in the strong sense). I need two modest assumptions. First, I suppose there are worlds at which the only natural external relations are spatiotemporal relations; worlds at which spacetime is empty will do. Second, I suppose that some such worlds are locally unified: individuals at the world that are neither adjacent nor overlapping are spatiotemporally separated; they (and their parts) are

³³ By ‘individual’, I always mean “thick particular.” See §2.6, which is the only place the distinction matters.

not related by any natural spatiotemporal relation (see §1.3). Now consider the sum of any two spatiotemporally separated individuals at any such locally unified world. Apply (PS) to that sum. In the possibility that results, the duplicate of the sum is composed of two individuals that are not only spatiotemporally *separated*, but spatiotemporally *isolated* as well; the intervening spacetime that unified them in the original world has been removed. By assumption, there are no non-spatiotemporal, natural external relations to unify them. So, they are *absolutely* isolated individuals: island universes in the strong sense.³⁴

2.5. *Lewis and (PS); Strong vs. Weak Denials of Necessary Connections.* The principle of solitude, as I see it, encapsulates an especially strong form of the Humean denial of necessary connections between distinct existents. Lewis cannot accept (PS) – not if I am right that it leads to the possibility of island universes. When Lewis champions the denial of necessary connections, then, it must be something weaker he has in mind. The difference between the strong and weak denial can be seen to arise from an ambiguity of quantifier scope. Let us say there are necessary connections between distinct existents if: there exists some x such that, necessarily, x coexists with some y distinct from x . This is ambiguous: the quantifier ‘some y ’ can be given wide scope or narrow scope. If given wide scope, we get:

(NC1) There exists some x and some y distinct from x such that, necessarily, x coexists with y .

When Lewis explicates the Humean denial of necessary connections as (in part) “anything can fail to coexist with anything else,” it is apparently the denial of (NC1) that he has in mind.³⁵ Indeed, that denial is all one needs to support the standard Humean arguments about laws and causation. If the quantifier ‘some y ’ is given narrow scope, we get (the equivalent of):

(NC2) There exists some x such that, necessarily, x does not exist by itself.

(PS) is the denial of (NC2).

I do not think a Humean can rest content with the denial of (NC1): that denial is too weak to capture the full plenitude of possibilities. To illustrate, consider a red ball and a green cube. (Pretend that colors and shapes are natural properties.) According to the denial of (NC1), it is possible for (a duplicate of) the red ball to exist without (any duplicate of) the green cube existing. But, surely, there is more. It is also possible for (a duplicate of) the red ball to exist without anything green existing. Yet this does not follow from the denial of (NC1). To allow for this possibility, we need a stronger principle of plenitude. But there is more. It is also possible for (a duplicate of) the red ball to exist without anything (else) colored, without anything (else) colored or extended. We have started down a slippery slope. (PS) – the denial of (NC2) – is waiting at

³⁴ I present this example, to a slightly different end, in Bricker (1993), p. 275.

³⁵ Lewis (1986), p. 88.

the bottom of that slope: it allows for all these possibilities in one fell swoop. There is no stopping short.³⁶

2.6. *Generalizations of (PS)*. There are two ways to generalize (PS) corresponding to two ways to unrestricted the quantifiers over worldbound individuals. We can quantify instead over *all* individuals, *transworld* individuals included (where a transworld individual is any sum of worldbound individuals from two or more worlds); and we can quantify, not only over *individuals*, or *thick* particulars, but over *thinned-down* particulars that have had some or all of their universals or tropes stripped away (if there are universals or tropes). Either generalization leads directly to the possibility of island universes (in the strong sense), without taking a detour through locally unified worlds.

Generalizing the first way gives this:

(GPS) *Generalized Principle of Solitude*. For any worldbound or transworld individual, possibly, a duplicate of that individual exists all by itself.

The possibility of island universes follows immediately from (GPS) by instantiating the quantifier to any transworld individual.

What motivates (GPS) is that there should be no restrictions on what can be actualized. On the realist position I espouse, actualization is primitive and absolute. What justification could there be for restricting actualizability to some parts of logical space – the worldbound individuals – while excluding it from others – the transworld individuals. Or look at it from God’s perspective. To restrict actualizability to worldbound individuals would be to put a limitation on God’s power to choose, one not grounded in any logical necessity. Suppose that, in surveying the worlds prior to actualization, God found that two or more worlds were tied for best. Why must God choose between actualizing one world, or the other? He’s all-powerful! He can simply say: “Actualize *those!*”

The second way of generalizing (PS) comes into play only for theories that accept universals or tropes. Any such theory will distinguish between “thick” particulars – what I have been calling “individuals” – and “thinned-down” particulars.³⁷ Thick particulars are maximal “bundles” of co-instantiated universals or tropes, perhaps together with a substratum, or “thin” particular. If one selectively strips off some universals or tropes from a thick particular, or from any of its thick particular parts, what is left is a thinned-down particular; if one strips off all these universals or

³⁶ There is another way to allow for the possibilities mentioned. One can instead invoke principles of plenitude for alien possibilities: instead of “removing” the world surrounding the red ball, one can “replace” it by something alien to actuality. But one would still be stuck with necessities involving the disjunction of *all* natural properties, actual and alien. On “alien” possibilities, see Lewis (1986), pp. 91-92.

³⁷ For example, see Armstrong (1989b), pp. 94-96.

tropes, and anything is left, what is left is a thin particular. Now, the generalization we seek simply applies unconditional actualization to all particulars:

(GPSP) *Generalized Principle of Solitude for Particulars*. For any particular (thick, thinned-down, or thin), possibly, a duplicate of that particular exists all by itself.

To illustrate, consider some actual individual – say, an electron. Among its natural properties, let us suppose, are having unit negative charge, having spin one-half, and having a mass of .511 MeV. Then, according to (GPSP), possibly, there exists a particle (an ordinary individual) just like an electron but with no charge property; or with no charge or spin property; or (on a substratum view) with no properties at all – a “bare” particular.

To get from (GPSP) to the possibility of island universes, one can start with any world containing two or more individuals and strip off relations, that is, polyadic universals or tropes. One way: apply (GPSP) to the world minus *all* polyadic universals or tropes. In the possibility that results, distinct individuals are absolutely isolated island universes. Another way is more selective, but works only for tropes: divide the world into two distinct individuals and apply (GPSP) to the world minus all polyadic tropes connecting the two individuals. In the possibility that results, the duplicates of the two individuals are absolutely isolated island universes.

That concludes my case for the possibility of island universes in the strong sense. My belief in this possibility is not an offhand modal opinion. It follows, in many ways, from general principles of plenitude applied to ordinary individuals and worlds, principles that are hard to deny for anyone with a broadly Humean approach to modality. It is not negotiable.³⁸

3. REALIST RESPONSES

3.1. *Primitive Worldmate*. Time for the realist to face the music. Island universes in the strong sense, I have argued, are metaphysically possible. That leaves the realist with two options: somehow revise the analysis of world so as to allow a world to have absolutely isolated parts; or somehow revise the standard analysis of modality so that island universes can be possible, without there being a world at which island universes exist. I begin with the first option.

According to Lewis, the alternative to rejecting the possibility of island universes is to posit a primitive worldmate relation.³⁹ A *world*, then, would be a region of logical space such that

³⁸ According to the classification of principles of plenitude in Bricker (1991), (PS) and its generalizations count as principles of plenitude for structures: they have implications as to which structures are possible (that is, possibly instantiated or actualized). (PS) implies that possibility is preserved under the operation of taking substructures; (GPS) implies, in addition, that possibility is preserved under the operation of taking sums of structures.

³⁹ Lewis (1986), pp. 71-72. For Lewis, rejecting island universes is “more credible” than positing primitive worldmate, though he doesn’t give his objections to the latter.

every part stands in this worldmate relation to every part, and to nothing else. But there is a major problem. A primitive worldmate relation is primitive modality, in at least two ways. First, the realist will need to posit necessary connections between the worldmate relation and other relations: necessarily, things are worldmates if they are spatiotemporally (or externally) related. Second, *general* facts of modality (facts expressible without reference to specific properties, relations, or things) will be made to depend upon how the worldmate relation is laid out in logical space; in particular, the very possibility under discussion, the possibility of island universes, will so depend.⁴⁰ To resort to a primitive worldmate relation, then, would spell defeat for the realist project of analyzing modality.

Even leaving the question of primitive modality to one side, it is far from clear how a primitive worldmate relation would solve the problem at hand. For what sort of relation is primitive worldmate? Not internal, of course: a duplicate of one of my worldmates need not be my worldmate. So it is external (since, presumably, being primitive, it is not extrinsic). But if it is a *natural* external relation, then it is no help at all! For then worldmates are externally connected, and no world has absolutely isolated parts, island universes in the strong sense.

Could primitive worldmate be some *non-natural* external relation? In that case, it would not genuinely unify, or tie together, its *relata*. (Compare non-identity, which is also, I claim, external but non-natural.) It would be best thought of, I think, not as a primitive relation, but as a plethora of primitive, non-natural properties, one per world; the one that applies to the actual world could do much of the work done by primitive actualization on my own theory.⁴¹

I cannot complain that these primitive, non-natural (and so non-qualitative) properties are mysterious, since I have helped myself to one for my own realist theory (though I can wonder at their abundance). But realism with primitive actualization has stark advantages over realism with primitive worldmate. It avoids positing primitive modality, and it upholds the ontological bifurcation, without which I claim realism is untenable, between the “actual” and the “merely possible.” Fortunately, given primitive actualization, we don’t need primitive worldmate to respond to the realist dilemma. There is a better way.

3.2. *The Amended Analysis of Modality: Class Version.* The realist has a second option: amend the standard analysis of possibility. It should be obvious by now how the emendation will go. If possibilities are represented not just by single worlds, but also by pluralities of worlds, then our problem is instantly solved: the possibility of island universes will be represented throughout

⁴⁰ Note that primitive actualization is not in the same boat: modal facts do not depend upon which region of logical space has been actualized; only truth so depends. Actuality is no more a modal notion than is truth.

⁴¹ I do not know whether Lewis would allow that worldmate could be both primitive and non-natural. In any case, since Lewis only considers primitive worldmate as a means to accommodate *spatiotemporally isolated* island universes, he only needs primitive worldmate to be non-spatiotemporal; it needn’t also be non-natural.

logical space, by every plurality of worlds.⁴² The solution comes in different versions, however, depending upon how “pluralities” are understood. If pluralities are understood as classes, we have:

Amended Analysis (Class Version): A proposition is (metaphysically) *possible* if and only if it is true at some (non-empty)⁴³ class of worlds.

The analysis requires that we make sense of truth *at a class of worlds*. The idea, of course, is to imagine that *all* the worlds in the class are actualized, and then to ask what would be true. More exactly, starting from a standard interpretation of a language in possible worlds semantics, we extend the interpretation as follows. The *domain* of a class of worlds is the union of the domains of the worlds in the class. An existentially or universally quantified sentence, when evaluated relative to a class of worlds, will have its lead quantifier restricted to the domain of that class. The extension of a predicate relative to a class of worlds is just the union of its extensions at the worlds in the class. That suffices to assign truth values to sentences of modal predicate logic.

It might seem rash to tinker with the standard analysis of possibility: it is a cornerstone of modern modal metaphysics, not only for realists, but for all who traffic in the language of possible worlds. Note, however, that the proposed emendation is conservative in its consequences. For one thing, although possibilities are added, no possibilities are taken away. Whatever was possible under the standard analysis – true at some world – remains possible under the amended analysis: true at the corresponding singleton world. For another thing, the added possibilities are extremely limited in scope. One new possibility, of course, is the proposition that there exist absolutely isolated individuals: it is true at any class of two or more worlds. And every other new possibility entails this proposition. So only propositions having to do with island universes change their possibility status under the amended analysis.

Possibility does not stand alone. There will have to be corresponding emendations in the analyses of other modal operators. Metaphysical necessity, of course, becomes truth at all *classes* of worlds. Restricted modal operators – nomological, doxastic, deontic, and so on – will have to be analyzed in terms of accessibility relations taking *classes* of worlds as their *relata*. For example, an agent’s (*de dicto*) beliefs will be represented, not by a class of doxastically accessible worlds, but by a class of *classes* of doxastically accessible worlds. Suppose an agent believes that island universes exist; then only classes of worlds containing two or more members are doxastically accessible (for that agent). Truth conditions for counterfactuals, too, will have to change: the closeness relation will take *classes* of worlds as its *relata*. To illustrate, consider

⁴² Warning: Lewis has also argued that possibilities need not be represented by possible worlds, but on entirely different grounds. Lewis claims that possibilities for an *individual* – possibilities *de re* – are represented, not by whole possible worlds, but by possible individuals, that is, parts of worlds. See Lewis (1986), pp. 227-235. We are concerned here, however, only with possibilities *de dicto*.

⁴³ On whether ‘non-empty’ should be dropped, see §4.2.

again a world that contains a single “wormhole” connecting two otherwise disconnected regions of spacetime. (This time make it deterministic, with a static spacetime.) Now consider the counterfactual: had there been no “wormhole,” there would have been absolutely isolated regions of spacetime. I take it this counterfactual is true. (Should one say instead that there would have been some other “wormhole”? That the antecedent is impossible?) What makes the counterfactual true is this: from the standpoint of the “wormhole” world, the closest class of worlds at which the antecedent is true is a doubleton, rather than a singleton; and the consequent is true at that closest class.

It would be onerous, to be sure, to have to rewrite the textbooks on possible world semantics, giving truth conditions relative to classes of worlds, instead of worlds. I recommend against it! Once one becomes convinced that it can be done, one need only indulge on those rare occasions – the present included – where the possibility of island universes comes into play.

3.3. *The Amended Analysis With and Without Absolute Actualization: Semantical Considerations.* Could Lewis accept the amended analysis so as to allow for the possibility of island universes? I think not. The amended analysis and Lewis’s “indexical” theory of actuality do not mix well. The problem comes, not with possibility, but with truth. Suppose I assert: “Island universes exist.” On the semantical framework that underlies the amended analysis, the truth or falsity of my utterance is to be evaluated relative to a *class* of worlds. But which class? The world at which my utterance occurs belongs to many classes of worlds, and without absolute actualization there is nothing to choose between them. I consider three options. (1) Stay as close as possible to the old method according to which the truth or falsity of an utterance is evaluated relative to the world at which the utterance occurs. On the new semantical framework, this becomes: my utterance, “island universes exist,” is true, *simpliciter*, if and only if it is true at the singleton whose sole member is the world at which my utterance occurs; otherwise, false, *simpliciter*. But, then, on semantical grounds alone, my utterance is false, *simpliciter*, since it is false at any singleton world. So, if we combine the Amended Analysis with option (1), we have it that my utterance is both contingently possible and *analytically* false. Not a happy combination. (2) Invoke supervaluations. Call any class of worlds containing the world at which my utterance occurs a *relevant* class. Then, my utterance, “island universes exist,” is true, *simpliciter*, if and only if it is true at all relevant classes; it is false, *simpliciter*, if and only if it is false (not true) at all relevant classes; otherwise, it is neither true nor false, *simpliciter*. But, then, on semantical grounds alone, my utterance is neither true nor false, *simpliciter*, since it is false at one relevant class (the singleton) and true at all the others. So, if we combine the Amended Analysis with option (2), we have it that my utterance is both contingently possible and *analytically* neither true nor false. Still not a happy combination. (3) Deny that utterances (at any rate, utterances that express contingent propositions) are ever true or false, *simpliciter*: there is only relative truth, truth relative to this or that class of worlds. But now I’ve lost my grip on the semantical enterprise. One job of an adequate semantics is to provide the machinery for moving from actual utterances to truth values. The assignment of truth *conditions* to utterances takes us only partway; there must also be a path from truth conditions to truth or falsity. Indeed, if there is no truth or falsity, *simpliciter*, in what sense are truth conditions of *truth*? Option (3), being semantically inadequate, is a non-starter. And that’s it, I think, for plausible options. The amended analysis, then, and the semantical framework that underlies it, are not for Lewis:

without absolute actualization, the alternative to rejecting the possibility of island universes, as Lewis said, is to posit primitive worldmate.⁴⁴

When the amended analysis is combined with absolute actualization, my assertion, “island universes exist,” acquires a definite truth value. One or more worlds has been actualized, and the truth or falsity of my assertion is evaluated relative to the class of actualized worlds. Of course, the assertion, “island universes exist,” is rather extraordinary. But it does no harm also to evaluate ordinary assertions relative to the class of actualized worlds. For most ordinary assertions, the quantifiers are implicitly restricted to the world at which the assertion occurs, or some part thereof; island universes, then, should they exist, would be irrelevant to the truth values of ordinary assertions. For some assertions, however, it is unclear whether or not island universes would be relevant. For example, when a physicist says, “nothing travels faster than light,” should her quantifiers be restricted to the world she inhabits, so that superluminal particles in other island universes would be irrelevant? Or should her quantifiers extend to other actualized worlds, so that her assertion would be falsified by an island universe at which light is not a first signal? I doubt there is any linguistic fact of the matter. There is no reason why the physicist should have bothered to decide which she means. Nor, then, should we decide. It is enough to note that we can allow for either interpretation.

So much for actual utterances. Should we also assign truth values to merely possible utterances? We can if we want. Typically when we ask whether a possible utterance is true, we are engaged in counterfactual thinking: *would* the utterance have been true, had it been made. In that case, we carry over information about actuality to the counterfactual situation. In particular, we carry over whether island universes do or do not exist, since, presumably, the existence of island universes is counterfactually independent of whether or not an utterance is made. A possible utterance of “island universes exist,” thus counterfactually considered, has the same truth value as an actual utterance of the same. We might mean something else, however, when we ask for the truth value of a possible utterance. We might be asking whether the possible utterance is true or false, *simpliciter*, true or false from its own perspective. In that case, the absolute facts about actuality are irrelevant. We are back to the three options considered above. (But this time, we needn’t worry that the options will be semantically inadequate, either by making truth values semantically determined where they should not be (options (1) and (2)), or by denying such truth values exist (option (3)); the assignment of truth values to *possible* utterances is not essential to the semantical enterprise, and can be introduced or jettisoned with impunity.) I think option (2), the method of supervaluations, gives the intuitively right results. When a merely possible person (speaking our language) says, “island universes exist,” what she says is neither true nor false, *simpliciter*. Although she intends her assertion to depend upon the facts of absolute actualization, there are no such facts (from her perspective), and so nothing to make her assertion either true or

⁴⁴ I have focused on the problem of evaluating the proposition expressed by an utterance. There is a parallel problem – with three parallel options – having to do with determining what proposition an utterance expresses in case the proposition expressed depends upon the world in which the utterance occurs (as happens, for example, with restricted modalities and counterfactuals).

false. But when the merely possible person says “I exist,” she speaks truly: her utterance is true at all relevant classes of worlds.

3.4. *The Amended Analysis: Plural Quantifier and Aggregate Versions.* On the amended analysis of modality, possibility is truth at some world, *or some plurality of worlds*. I began with a class version of the analysis – interpreting pluralities as classes – not because I favor that version, but because among contemporary metaphysicians classes are familiar and, for the most part, accepted tools of the trade. In this section, I consider two further versions, and state my preferences. Because the philosophical arguments that favor one version over another largely cut across the issues of this paper, I will be brief. All three versions allow equally for the possibility of island universes.

The version I favor analyzes modality as plural quantification over worlds, and plural quantification, I have been convinced, is not to be reduced to singular quantification over classes (or class-like entities).⁴⁵ Consider first an ordinary language example. I am deciding which books to put on a shelf. You warn: “Some books will bring down the shelf.” You deny, however, that any single book will bring it down. Then you have quantified plurally over books. In longwinded paraphrase: there are some books such that *they* will bring down the shelf. You have not thereby quantified over anything other than books; in particular, you have not unwittingly, surreptitiously also quantified over, or somehow trafficked in, classes (or class-like entities). Just as the predicate ‘will bring down the shelf’ may be either singular or plural, I propose to interpret the relational predicate ‘is true at’ as either singular or plural in its second argument place, allowing it to take either singular or plural quantifiers over worlds. Thus, a proposition may be true at some world, or true at some worlds, where the latter does not entail the former. Recasting the amended analysis in terms of plural quantification, we have:

Amended Analysis (Plural Quantifier Version). A proposition is (metaphysically) *possible* if and only if it is true at some world, or some worlds.

Of course, a proposition is true at some worlds just in case it is true at the class containing those worlds as members according to the account given in §3.2. So the class version and the plural quantifier version do not differ as to what propositions are possible; in particular, on both versions, it is possible that island universes exist.

Why, if the versions are extensionally equivalent, do I prefer the plural quantifier version? Because, first, on the class version, one cannot say that a proposition is possible without being ontologically committed to classes; and that is plainly wrong. A (Harvard) nominalist who refused to countenance classes could nonetheless consistently aver that island universes are possible. Classes have nothing to do with it. Second, the plural quantifier version deviates less from the standard analysis, and so is less suspect. What is central to the standard analysis is that modal operators are quantifiers over worlds. The plural quantifier version of the amended

⁴⁵ For convincing arguments, see Boolos (1984), Boolos (1985), and Lewis (1991). For interesting dissent, see Resnik (1988) and Hazen (1993).

analysis doesn't reject that. It just adds: they may be plural, as well as singular, quantifiers. A final reason I postpone until §4.2.⁴⁶

A third version of the amended analysis interprets truth at a plurality of worlds, not in terms of classes or in terms of plural quantification, but in terms of *aggregates* (that is, mereological sums). Thus,

Amended Analysis (Aggregate Version). A proposition is (metaphysically) *possible* if and only if it is true at some aggregate of worlds.

Since worlds do not overlap, aggregates of worlds and (non-empty) classes of worlds are in one-to-one correspondence. That guarantees that the aggregate version is extensionally equivalent to the others.

The aggregate version and the plural quantifier version share a common advantage: assertions of possibility do not carry ontological commitment to classes. Of course, the aggregate version is committed to aggregates of worlds. But, appearances notwithstanding, that is no disadvantage: aggregates, I have been convinced, are an ontological free lunch; if one is committed to some things, then one is committed to the aggregate of those things, willy-nilly.⁴⁷ Why, then, do I prefer the plural quantifier version to the aggregate version? I have three reasons. First, as before, the plural quantifier version deviates less from the standard analysis by analyzing possibility as a quantifier over worlds. Second, if one accepts the aggregate version, one is tempted to generalize the analysis, so that a proposition is possible also if it is true at some *part* of some world. But that would be wrong. It would make possible the proposition: something exists that is spatiotemporally related to something that doesn't exist. And then we might wonder whether, if we went far enough out in spacetime, we would encounter the merely possible! That's absurd. The third reason I postpone until §4.2.

3.5. *Aggregates of Worlds are Worlds: A Third Realist Response?* The realist, I have said, can allow for the possibility of island universes in either of two ways: revise the criterion of demarcation for worlds, or amend the standard analysis of modality. With a mere shift of terminology, however, a solution of the second sort – the aggregate version of the amended analysis – can be transformed into a solution of the first sort. Say that a *world*, under a revised criterion, is any aggregate of worlds, under the original criterion. Then, the standard analysis applied to worlds in the revised sense is identical with the aggregate version of the amended analysis applied to worlds in the original sense; the two approaches are essentially equivalent. Since the difference is only terminological, there can be no substantive reason to prefer one

⁴⁶ Note that one may reject the class version but still choose to formulate possible world semantics in terms of classes of worlds (if, that is, one believes in classes). The ontological commitments of the metalanguage within which we do semantics for natural language may transcend the ontological commitments of natural language.

⁴⁷ See Lewis (1991).

approach to the other. We do not have a third realist solution. The question can only be: which way of speaking deviates less from established usage?

First off, the word ‘world’ as it occurs in modal metaphysics is a philosophical term of art, not a part of ordinary language. Moreover, since the standard analysis of modality is couched in terms of worlds, it can be no more a part of ordinary language than is the notion of world itself. To answer our question, then, we must look to established philosophical usage. For realists from Leibniz through Lewis, worlds have been essentially unified, and they have either not overlapped, or not overlapped extensively. The revised criterion of demarcation would be a radical departure from that tradition.⁴⁸ On the other hand, the amended analysis of modality – at least, the plural quantifier version – preserves the central core of the standard analysis: modal operators are quantifiers over possible worlds. It counts as a minor modification. Thus, I recommend on terminological grounds amending the standard analysis.⁴⁹

4. FURTHER APPLICATIONS

Thus far, I have dealt with one problem for Lewis’s modal realism, the problem of island universes. I have argued that, if a realist accepts the amended analysis of modality (and absolute actualization), the problem is easily resolved. Three further problems for Lewis’s modal realism are likewise easily resolved: (1) an unqualified principle of compossibility can be accepted, thereby allowing for the possibility of universal actualization; (2) the possibility of nothing can be endorsed, if desired, with an appropriate modification of the amended analysis; and (3) the principle of the identity of indiscernible worlds, undecidable on Lewis’s theory, can be decisively refuted. I treat these three problems in turn, followed by a brief conclusion.

4.1. *Universal Actualization and Lewis’s Principle of Recombination.* It is natural to think that only part of logical space has been actualized: flying pigs, planets of pure gold, these are merely possible beings existing nowhere in actuality. But is it not at least conceivable, and metaphysically possible, that *all* of logical space has been actualized, that everything possible exists? It follows from principles of plenitude that I accept – for example, from (GPS) – that the answer is “yes.”

⁴⁸ However, non-realists who gloss ‘possible world’ as ‘counterfactual situation’ – such as Kripke (1980) – might find (the non-realist analogue of) the revised criterion compatible with their usage.

⁴⁹ Ted Sider suggested in conversation (in 1992) the idea that every possible individual is a world, transworld individuals included; that aggregates of worlds are worlds is a special case. It wasn’t until I worked out the ideas of this paper the following year that I came to appreciate the insight behind his suggestion. When this work was presented at Princeton in March, 1996, I was shown an unpublished manuscript by Richard B. Miller, “Actuality, Island Universes and Schrodinger’s Cat,” in which the idea that aggregates of worlds are worlds is adopted to allow for the possibility of island universes.

Indeed, many philosophers, from ancient times to modern, have defended a principle of plenitude for actuality according to which whatever can exist, does.⁵⁰ Their reasoning, if set within the present Humean approach to possible worlds, would go like this: God, being perfectly good, will choose to actualize the best; but more is always better than less; so, God will choose to actualize, not this world or that, not these worlds or those, but all the worlds in logical space. I won't vouch for the theology; but the possibility of universal actualization seems perfectly coherent.

The amended analysis allows for the possibility of universal actualization: 'everything possible exists' is true at the class (aggregate, plurality) of all possible worlds. On the standard analysis, however, universal actualization, literally interpreted, is out of the question: 'everything possible exists' is true at no world. Is some non-literal interpretation of universal actualization compatible with the standard analysis? Lewis accepts a "principle of recombination" which, *if left unqualified*, entails that anything can coexist with anything, or, more generally, that, for any things, those things can coexist. Since the "things" may be in different worlds, possible coexistence is to be understood in terms of duplicates: for any things, some world contains distinct duplicates of those things. The *unqualified* principle of recombination, then, would allow for the possibility of universal actualization in an attenuated sense: possibly, every possible intrinsic nature is (distinctly) instantiated.

But a well-known argument due to Forrest and Armstrong, if appropriately beefed up, shows that the unqualified principle of recombination leads to contradiction: the big world that contains distinct duplicates of all the worlds would, in a sense that can be made precise, have to be bigger than itself.⁵¹ To avoid the contradiction, Lewis adds a qualifying proviso: for any things, those things can coexist *size and shape of possible spacetimes permitting*. When it comes to all things, no possible spacetime will be big enough. That avoids the contradiction, but at a substantial cost: it rejects as impossible what a great many philosophers throughout history have thought possible, even actual.

When we switch to the amended analysis, the Forrest-Armstrong argument loses its bite. Although no one world mirrors all the worlds in logical space, that no longer rules out the possibility of universal actualization. (GPS) can serve as a pure, unqualified Humean principle of plenitude for compossibility: *contra* Leibniz, all things are compossible. Precisely what recombination principle to adopt *in addition to* (GPS) is a matter for another occasion.

⁵⁰ See Lovejoy (1936) for a detailed account of the history of the "principle of plenitude" (for actuality), which holds in part "... that no genuine potentiality of being can remain unfulfilled, that the extent and abundance of the creation must be as great as the possibility of existence ..." (p. 52).

⁵¹ The argument is in Forrest and Armstrong (1984), and reformulated by Lewis in Lewis (1986). As pointed out in Nolan (1996), the contradiction does not follow from any premises explicitly presented either by Forrest and Armstrong, or by Lewis. The gap can be filled, however, by appealing to either a principle of plenitude for possible structures, or a principle of plenitude for alien natural properties. I hope to elaborate elsewhere.

4.2. *Nothing*. I believe that there might have been nothing at all. Nothing *physical*, that is; not even empty spacetime. Logical space would have existed as unrealized potentiality, waiting upon an absent creator. This belief is controversial, to be sure – more so, I think, than the possibility of island universes or universal actualization. I defend it, as best I can, at the end of this section.

Realism about possible worlds when combined with the standard analysis of modality cannot accommodate the possibility of nothing. Any part of any world exists at that world, and any world has itself as a part; so there is no world at which nothing exists. Then, given the standard analysis, it is impossible that nothing exist.⁵² Nor does switching to the amended analysis, as it stands, help. To accommodate the possibility of nothing, the amended analysis must be modified so as to include, in effect, a “null plurality” of worlds. But now it matters which version of the amended analysis the realist accepts; for not all versions can incorporate the modification in a natural way.

Consider first the aggregate version. Here we would have to allow a “null aggregate” of worlds, and say: possibility is truth at some aggregate of worlds, the null aggregate included. But there is no such thing as a null aggregate! On this version, the possibility of nothing is really an *ad hoc* special case; it does not follow from the analysis of possibility as a quantifier over aggregates of worlds.

Consider next the class version. Here the modification seems to be better off: possibility is truth at some class of worlds, the null class included. But is the null class ontologically more respectable than the null aggregate? I doubt it.⁵³ In which case, the possibility of nothing is just as *ad hoc* on the modified class version, as on the modified aggregate version.

Consider, finally, the plural quantifier version. Here the modification seems to be in trouble: plural quantifiers in English do not range over things “in the null way” that would be required for the modification. But the trouble is one of ordinary language, not logic or metaphysics, one of expression, not understanding. For we understand second-order logic with the second-order monadic quantifiers ranging over all subclasses of the domain, the null class included. And we understand, I have claimed, how to interpret quantification over non-empty classes as ontologically innocent plural quantification. To hold that ontological commitment to non-empty classes can be eliminated in this way, but not ontological commitment to the null class, would be absurd! The fact that the quantifiers of second-order logic do not match up neatly with the plural quantifiers of ordinary language is a mere technicality – of no more importance to logic or metaphysics than the fact that the quantifiers of first-order logic do not match up neatly with the singular quantifiers of ordinary language.

⁵² See Lewis (1986), pp. 73-74.

⁵³ On Lewis’s mereological theory of classes – classes are aggregates of singletons – the null class and the null aggregate are in exactly the same boat (though one might choose to introduce something arbitrarily to play the theoretical role attributed to the null class, say, in mathematics or possible-world semantics). See Lewis (1991), pp. 10-15.

That leaves a small problem of expression. If we want to express the modified amended analysis in something resembling English, we will need to coin a phrase. Say that a proposition is *true at nothing* just in case, intuitively, had no world been actualized, the proposition would have been true. To be more precise, we can let some developed version of free logic – logic over an empty domain – be our guide. The plural quantifier version of the amended analysis then becomes: possibility is truth at some world, or at some worlds, or at nothing. Since the proposition that nothing exists is true at nothing, it comes out possible, as desired, even though it is true at no world. Note that being true at nothing is not the same as being true at no world: a contradiction, for example, though true at no world, is not true at nothing; if it were, contradictions would be possible.

On the surface, the modified amended analysis looks like an *ad hoc* collection of clauses. But deeper down, its content is seamless. When asked which worlds might be actualized, we answer: all, or some, or none. We cover the full range of quantifiers. What would be arbitrary would be to leave off the ‘none’.

The possibility of nothing does not follow from the Humean principles of plenitude accepted in §2.4 and §2.6. They need to be strengthened. That is easily done. The relevant question is then: do the arguments that served to motivate the original principles also serve to motivate the strengthened versions? Consider first (PS). Say that a worldbound individual x at a world w is *strongly contingent* if and only if, possibly, x fails to exist without anything taking its place; that is to say, possibly, all and only duplicates of the parts of w that do not overlap x exist. Then (PS) is equivalent to: any worldbound individual *other than a whole world* is strongly contingent. To strengthen (PS), we simply omit the italicized clause: any worldbound individual is strongly contingent. The possibility of nothing now follows by applying the strengthened (PS) to whole worlds.

My defense of (PS) rested upon the claim that actualization is unconditional: whether or not some thing can be actualized does not depend upon whether or not anything else is actualized along with it. To defend strengthened (PS), we need instead a form of unconditional *de*-actualization: whether or not some thing can be *de*-actualized – can fail to exist without anything taking its place – does not depend upon whether or not anything else exists at its world. For consider some worldbound individual x that is not a whole world. By ordinary (PS), there is a world v that is a duplicate of x . If we do *not* strengthen (PS), then x is strongly contingent (can be *de*-actualized) but v is not (cannot), even though they differ only extrinsically: v , but not x , exists all by itself. I accept unconditional *de*-actualization; but I do not think it should be allotted a fundamental role in a theory of plenitude alongside unconditional actualization. A case for the possibility of nothing based upon unconditional *de*-actualization is weak.

A stronger case can be made by considering the argument behind (GPS). To strengthen (GPS), we simply move from an aggregate to a plural quantifier formulation, allowing plural quantifiers, as above, to range over “nothing”: for any things, possibly, distinct duplicates of those things exist all by themselves. My defense of original (GPS) rested upon the claim that there should be no restrictions on what can be actualized. The argument applies no less to strengthened (GPS). Consider a do-nothing God, content to contemplate the eternal verities. To hold that He *must* actualize some world, or some worlds, is to restrict His power to choose:

being all-powerful, He can choose to actualize nothing. And, if rejecting the possibility of nothing would be a restriction on God's power to choose, so too would it be a restriction on primitive actualization, a restriction on what can be actualized.

4.3. *Indiscernible Worlds*. Any realist account of possible worlds must face the question whether distinct worlds are ever qualitative duplicates of one another, or, equivalently, given the absolute isolation of worlds, the question whether there are counterexamples to the identity of qualitatively indiscernible worlds.⁵⁴ For if worlds are particulars – as opposed, say, to properties or universals – then duplication is not ruled out categorically as incoherent. David Lewis writes: “My modal realism does not say whether or not there are indiscernible worlds; and I can think of no very weighty reason in favor of one answer or the other.”⁵⁵ Indeed, as long as we hold to the standard analysis according to which possibility is truth at a single world, modal intuitions are powerless to decide the issue: no judgment of possibility would be affected by the presence or absence of indiscernible worlds. There is a mismatch between the framework of possible worlds and the judgments of possibility that the framework serves to interpret. A feature of the framework, the existence or non-existence of indiscernible worlds, appears arbitrary and artificial – an ontological dangler, if you will. Something, I think, is amiss.

The amended analysis of modality sets this right. It allows familiar arguments against the identity of indiscernible worldmates to be straightforwardly applied to the worlds themselves. For, surely, if island universes are possible, then it is possible for the islands to be qualitatively very similar; and, if very similar, why not exactly alike?⁵⁶ But the possibility of duplicate island universes, on the amended analysis, requires the existence of distinct but indiscernible worlds. The identity of indiscernible worlds, then, is false.

Is this argument conclusive? It does assume that the possibility of indiscernible island universes is to be analyzed in a way analogous to the possibility of *almost* indiscernible island universes: since the latter possibility is made true by a pair of distinct worlds, so is the former. But this assumption cannot plausibly be challenged by a realist. For on what grounds would the possibilities be treated differently? Granted, the possibility of indiscernible island universes is made true by a single *type* of world, instantiated twice-over, whereas the possibility of almost indiscernible island universes is made true by two distinct types. But to analyze possibility in terms of *types* of world is to move away from realism, as here characterized, and to identify worlds instead with uninstantiated properties or universals.⁵⁷ As long as worlds are taken to be particulars, the argument, I think, is conclusive.

⁵⁴ On duplicates vs. indiscernibles, see Lewis (1986), pp. 62-63.

⁵⁵ Lewis (1986), p. 157.

⁵⁶ This is an adaptation of Robert Adams's “argument from the possibility of almost indiscernible twins.” See Adams (1979), pp. 17-19.

⁵⁷ As is done, for example, in Forrest (1986) and Stalnaker (1976).

4.4. *Conclusion.* The standard analysis of modal operators as individual quantifiers over worlds is well-entrenched, but not sacrosanct. Analyzing modal operators instead as plural quantifiers over worlds (or individual quantifiers over classes, or aggregates, of worlds) has a lot to recommend it for the realist. No ordinary possibility judgments are affected by the shift, and the newly added possibilities are theoretically very satisfying: they allow the realist to accept the plenitude of possibilities to its fullest extent. That includes the possibility of island universes (in the strong sense), even the ultimate possibility of universal actualization. It includes (on the plural quantifier version) the possibility of nothing. And the amended analysis resolves the otherwise mysteriously aloof identity of indiscernible worlds. The price for all this is absolute actualization. But, if I'm right, that price is not so great as is often supposed.

NOTES

It is a privilege and a pleasure to contribute to a volume on the philosophy of David Lewis. His work has been a fountain of philosophical inspiration – and good sense – wherein I continually replenish myself. Portions of this paper were presented at Princeton University in March, 1996. Thanks to David Lewis and Ted Sider for helpful discussion over a number of years.

REFERENCES

- Adams, Robert (1979): "Primitive Thisness and Primitive Identity," *The Journal of Philosophy* 76, pp. 5-26.
- Armstrong, D. M. (1989a): *A Combinatorial Theory of Possibility*, Cambridge, UK: Cambridge University Press.
- Armstrong, D. M. (1989b): *Universals: An Opinionated Introduction*, Boulder, Colorado: Westview Press.
- Bigelow, John and Pargetter, Robert (1987): "Beyond the Blank Stare," *Theoria* 53, pp. 97-114.
- Boolos, George (1984): "To Be is to Be the Value of a Variable (or to Be Some Values of Some Variables)," *Journal of Philosophy* 81, pp. 430-449.
- Boolos, George (1985): "Nominalist Platonism," *The Philosophical Review* 94, pp. 327-344.
- Bradley, F. H. (1893): *Appearance and Reality*, London: George Allen and Unwin.
- Bricker, Phillip (1991): "Plenitude of Possible Structures," *The Journal of Philosophy* 88, pp. 607-619.
- Bricker, Phillip (1993): "The Fabric of Space: Intrinsic vs. Extrinsic Distance Relations," in P. A. French, T. E. Uehling, Jr., and H. K. Wettstein (eds.), *Midwest Studies in Philosophy*, Vol. 18, Notre Dame, Indiana: University of Notre Dame Press, pp. 271-294.
- Bricker, Phillip (1996): "Isolation and Unification: The Realist Analysis of Possible Worlds," *Philosophical Studies* 84, pp. 225-238.
- Forrest, Peter (1986): "Ways Worlds Could Be," *Australasian Journal of Philosophy* 64, pp. 15-24.
- Forrest, Peter and Armstrong, D. M. (1984): "An Argument Against David Lewis' Theory of Possible Worlds," *Australasian Journal of Philosophy* 62, pp. 164-168.
- Hazen, A. P. (1993): "Against Pluralism," *Australasian Journal of Philosophy* 71, pp. 132-144.
- Kant, Immanuel (1781): *Critique of Pure Reason*, translated by Norman Kemp Smith, London: MacMillan.
- Kripke, Saul (1980): *Naming and Necessity*, Cambridge, Mass.: Harvard University Press.
- Leslie, John (1989): *Universes*, London: Routledge.
- Lewis, David (1970): "Anselm and Actuality," *Noûs* 4, pp. 175-188.
- Lewis, David (1973): *Counterfactuals*, Cambridge, Mass.: Harvard University Press.

- Lewis, David (1986): *On the Plurality of Worlds*, Oxford: Blackwell.
- Lewis, David (1991): *Parts of Classes*, Oxford: Blackwell.
- Lewis, David (1992): "Critical Notice: Armstrong, D. M., *A Combinatorial Theory of Possibility*," *Australasian Journal of Philosophy* 70.
- Lovejoy, Arthur (1936): *The Great Chain of Being*, Cambridge, Mass.: Harvard University Press.
- Miller, Richard B. (unpublished): "Actuality, Island Universes and Schrödinger's Cat."
- Nolan, Daniel (1996): "Recombination Unbound," *Philosophical Studies* 84, pp. 239-262.
- Resnik, Michael (1988): "Second-Order Logic Still Wild," *The Journal of Philosophy* 85, pp. 75-87.
- Russell, Bertrand (1950): *An Inquiry into Meaning and Truth*, London: George Allen and Unwin.
- Skyrms, Brian (1976): "Possible Worlds, Physics and Metaphysics," *Philosophical Studies* 30, pp. 323-332.
- Stalnaker, Robert (1976): "Possible Worlds," *Noûs* 10, pp. 65-75.
- Strawson, P. F. (1966): *The Bounds of Sense*, London: Methuen.
- Weingard, Robert (1976): "On the Unity of Space," *Philosophical Studies* 29, pp. 15-20.