PHILOSOPHICAL TOPICS
VOL. 26, NO. 1 & 2 SPRING & FALL 1999

Intrinsicness[1]

S. Yablo
MIT

I. INTRODUCTION

You know what an intrinsic property is: it's a property that a thing has (or lacks) regardless of what may be going on outside of itself. To be intrinsic is to possess the second-order feature of stability-under-variation-in-the-outside-world.

You probably know too why this is hopeless as a philosophical account of intrinsicness. "Variation in the outside world" has got to mean "variation in what the outside world is like *intrinsically*." Otherwise every property G is extrinsic; for a thing cannot be G unless the objects outside of it are *accompanied* by a G.[2]

But, although the naive account is circular, it is not beyond salvage. Leave aside for a minute the problem of saying what *in general* constitutes intrinsic variation in the outside world. A special case of the notion can be defined independently. This special case, suitably iterated, turns out to be enough; suitably iterated the special case *is* the general case.

II. ANALYSIS

Before introducing the special case, a word about motivation. A philosophical account of X ought if possible to bring out de jure relations between X

and other notions. It is not enough if X covaries as a matter of fact with ABC, even if the de facto covariation is counterfactually robust. (I include here metaphysically necessary covariation as the limiting case.) How well has this condition—the "de jure condition," let's call it—on a successful analysis been respected in the literature on intrinsicness?

The first proposal worth mentioning is due to Jaegwon Kim. G is intrinsic according to Kim iff it is compatible with *loneliness:* the property of being unaccompanied by any (wholly distinct, contingently existing) thing.[3] It seems of the essence of intrinsicness that an intrinsic property should be possessable in the absence of other things; and so Kim's account does fairly well with the de jure condition. The problem with the account is that, as Lewis noticed back in 1983, it gives the wrong results.[4] Loneliness is as extrinsic as anything; but since it is a property compatible with loneliness (!), Kim would have to call it intrinsic.

What about Lewis's own account, according to which G is intrinsic iff given any x and y with the same natural properties, x is G iff y is G?[5] This gives more accurate results, but an element of de facto-ness has now intruded. If some natural property H should fail to be intrinsic, the account will overgenerate; it will call H "intrinsic" regardless. You may say that such a situation will never arise, and you may be right. The problem is that its never arising seems no more than a lucky accident.[6] There is no *in principle* reason why theorists of the quantum domain should not find themselves forced by nonlocality phenomena to count certain extrinsic properties as "ground floor" and natural. Because there is nothing in the nature of intrinsicness to prevent this, it is a matter of luck that Lewis's story succeeds as well it does.

A modified account developed with Rae Langton draws only on a particular aspect of naturalness, viz., nondisjunctiveness.[7] According to Langton and Lewis, G is "basic intrinsic" iff (i) G and its negation are independent of loneliness and accompaniment and (ii) G and its negation are nondisjunctive. The intrinsic properties are the ones that never distinguish between things with the same basic intrinsic properties.

If you think that nondisjunctiveness comes closer than naturalness to being in the same conceptual ballpark as intrinsicness, then the modified account is more de jure than its predecessor. But an element of de facto-ness remains. Consider the property R of being the one and only round thing. R satisfies (i) and the first half of (ii); if it avoids being (basic) intrinsic, then, that is because it is the negation of the disjunctive property S of being non-round or else round and accompanied by a round thing.

Is S disjunctive, though? That it can be expressed in disjunctive terms doesn't count for much; roundness too can be expressed that way, e.g., as the disjunction of round-and-charged with round-and-not-charged. Better evidence of disjunctiveness would be a finding that S's "disjuncts"—the property of not being round, and that of being round and accompanied by something round—were perfectly natural. But they clearly are not, and so Langton and Lewis are led to maintain that S is at least much *less* natural than its "disjuncts." This means in particular that

> S = accompanied by something round *if* round oneself

is much less natural than

> T = accompanied by something round *and* round oneself.

It may well be. But there is something uncomfortable about taking an intrinsicness-fact that is very clearcut—the fact that R is not intrinsic—and putting it at the mercy of something as controversial, and (apparently) irrelevant, as the relative naturalness of S and T. One feels that it ought not to *matter* to the issue of R's intrinsicness where S and T come out on the naturalness scale.

Of course, we don't want to be in the position of asking too much. Who is to say that intrinsicness bears *any* worthwhile de jure relations to other notions? Or maybe it bears them only to notions less fundamental than itself—notions that ought to be explained in terms of intrinsicness rather than the other way around.

The answer to this is that intrinsicness *does* appear to line up in nonaccidental ways with something quite fundamental: the relation of part to whole. It seems, for instance, as de jure as anything that

> if u is part of v, then u cannot change intrinsically without v changing intrinsically as well.

And that

> if u is part of v, then u and v have a region of intrinsic match.

And that

> if u is properly part of v, then u and v have intrinsic differences.

So the materials for a more principled account are not obviously lacking. What remains to be seen is whether we can parlay one-way conditionals like the above into a de jure biconditional with intrinsicness all by itself on the left hand side.

III. EXPANSION

The last of our three conditionals says that if one thing is properly part of another, then the two are intrinsically unalike. If you believe that, then

you'll agree with me that *one* good way to arrange for intrinsic variation in the world outside of x is to *expand it:* add on something new. This immediately gives a necessary condition on intrinsicness:

(*) G is intrinsic only if necessarily, whether a thing is or is not G cannot be changed by adding a part to its containing world.

I say that the necessary condition is also sufficient. A counterexample would be a G that could be made to vary through intrinsic variation in the outside world, but not through the particular *sort* of intrinsic variation contemplated here: expansion. But for reasons about to be explained, such a G is not possible.

IV. WHY EXPANSION IS ENOUGH

Remember the intuition that we started with: If G is an extrinsic property, then it can be lost (gained) through intrinsic variation in the outside world. It follows from this that there are w and w' such that x has G in w but not w', and the changes rung on w to get w' are *outside* of x.

Because the part of w that lies within x is left absolutely untouched, w and w' have a part in common: a part extending at least to x's boundaries and possibly beyond. Focus now on this shared part. It could have existed all by itself; why not? Another way to put it is that there is a self-standing world (call it w'') consisting of the shared part and nothing else.

The question is, does G obtain in w'' or does it not? However you answer, there are worlds u and v such that (i) x is part of u is part of v, and (ii) G is either lost or gained as you move from u to v. Suppose first that G does obtain in w''; then G is lost as we move from $u = w''$ to $v = w'$. If, on the other hand, G doesn't obtain in w'', then G is gained as we move from $u = w''$ to $v = w$. Either way, G is sensitive to positive mereological variation in the outside world. Condition (*) is not just necessary for intrinsicness but a sufficient condition as well.

V. ASSUMPTIONS

A few implicit assumptions must now be brought to light. Or let me stress: *most of them are for simplicity only* and will eventually be given up. This applies in particular to assumptions (1)–(3), which together go under the name of "modal realism." Later on (see the appendix), we'll be abandoning possible worlds altogether. But for the time being our outlook is very close to that of David Lewis. We accept:

(1) *Pluralism:* This actual world is only one of a large number of possible worlds. It is possible that BLAH iff in at least of these worlds, BLAH is the case.
(2) *Possibilism:* These other worlds exist—they are part of what there is—but they are no part of actuality.
(3) *Concretism:* Other possible worlds are entities of the same basic sort as this world. They are maximal concrete objects that just happen not to be the same maximal concrete object that we around here inhabit.
(4) *Mereologism:* There is only one part/whole relation worth taking seriously, a relation answering to the axioms of mereology.
(5) *Recombinationism:* The space of worlds is closed under arbitrarily reshuffling of world-parts. As Lewis puts it, "patching together parts of different possible worlds yields another possible world."[8]

Finally, a seeming corollary of (5),

(6) *Inclusionism:* Some worlds contain others as proper parts.

I said that our outlook was "close to" that of Lewis; I might have said that we are going to be more Lewisian than Lewis himself. For it turns out that Recombinationism is *not,* the quoted passage notwithstanding, a view that Lewis accepts. ("An attached head does not reappear as a separated head in some other world, because it does not reappear at all in any other world."[9]) Nor does Lewis accept Inclusionism. The reason is the same in both cases. Both assumptions represent worlds as having parts in common; and the sharing of parts is, according to Lewis, strictly forbidden. (A section of *Plurality of Worlds* is called "Against Overlap." What part of "Against" don't you understand?)

VI. ACCIDENTAL INTRINSICS

If one looks, however, at what Lewis actually says in "Against Overlap," his position turns out to be not quite as adamant as suggested:

> my main problem is not with the overlap itself. Things do have shared parts, as in the case of the Siamese twins' hand . . . what I do find problematic—inconsistent, not to mince words—is the way the common part of the two worlds is supposed to have different properties in the one world and in the other.[10]

Reading on, we learn what the inconsistency is. According to friends of overlap,

> Humphrey, who is part of this world and here has five fingers on the left hand, is also part of some other world and there has six fingers on his left hand. . . . He himself—one and the same and

> altogether self-identical—has five fingers on the left hand, and he has not five but six. How can this be? You might as well say that the shared hand of the Siamese twins has five fingers as Ted's left hand, but it has six fingers as Ned's right hand! That is double-talk and contradiction. Here is the hand. Never mind what else it is part of. How many fingers does it have?[11]

Lewis calls this the *problem of accidental intrinsics.* What matters to us is that it is a problem raised by some cases of overlap, but not all.

The problem does not arise, Lewis says, for Humphrey's *essential* properties,[12] however intrinsic they may be. No explanation is needed of how Humphrey can have different essential properties in different worlds, because he doesn't.

"Neither [does] it arise for Humphrey's extrinsic properties, however accidental." A thing's extrinsic properties are, implicitly, relations it bears to its surroundings. And there is no contradiction whatever in bearing a relation to one set of surroundings (being shorter than anyone in those surroundings, say) that you do not bear to another.

VII. NOT ALL OVERLAP IS BAD

As far as the main argument of "Against Overlap" goes, a part of one world *can* recur in another world *provided that its intrinsic properties do not vary between the two worlds.*

One place we might seek assurances on this score is in the nature of the shared item. If the item was of a sort that possessed its intrinsic properties *essentially,* there would be no danger of its changing in intrinsic respects between the one world and the other.

A second place assurances might be sought is in the nature of the *relation* between the two worlds. Even if the item did not retain its intrinsic properties across *all* worlds, it might stay intrinsically the same across certain particular *pairs* of worlds—such as, for example, those standing in the part/whole relation. If so, there would be no objection to saying that the item was shared by those particular pairs of worlds.

And now our cup overfloweth; for our interest here is in a *case* of overlap—the case where one of two worlds is included in the other—that eludes Lewis's strictures in *both* of the ways just mentioned.

It eludes them in the first way because worlds are generally understood to have all of their intrinsic properties essentially.[13] World-inclusion avoids the problem of accidental intrinsics because it is *world*-inclusion.

Also, though, it avoids the problem because it is world-*inclusion.* One world is not part of another unless it recurs in it with the same intrinsic properties it had qua self-standing world.[14] Example: a world just like this one up to five minutes ago, when irresponsible atom-smashing experiments occurred causing everything to pop out of existence *now,*[15] is *not* an initial segment of actuality—whereas a world just like this one up to now, when everything *miraculously* pops out of existence, *is* an initial segment of actuality.[16]

VIII. DEFINITION

I said that G is intrinsic iff, supposing that x has (lacks) G in w, G cannot be canceled (introduced) by moving to w', where w' is w with the addition of a new part.

This may seem unnecessarily cagey. Why "canceled"? Why not just say that there can be no w' of the indicated type in which G fails to be exemplified by the original object x: the object that *did* exemplify G back in w?

One worry would be that x might not *persist* into w'. But if we stick to our policy of departing from Lewis's metaphysics or only when absolutely necessary, this eventuality can be ruled out.

Remember, Lewis is a *mereologist:* he maintains that there is only one containment relation worth taking seriously, a relation (partially) characterized by the axioms of mereology.[17]

How does mereologism eliminate the danger of x not surviving the trip to w'? The story so far is that w and w' have a part in common that contains x. If containment is mereological part/whole, this means that x is a mereological part of the mereological overlap between w and w'. But the overlap between w and w' is (trivially) part of w'. So we can conclude by the transitivity of part/whole that x is part of w', i.e., that w' contains x.

Now, if x persists into w', then it seems clear that x is the thing in w' that G had better still attach to, if G wants to be regarded as intrinsic.[18] This gives us our first official definition of intrinsicness ('<' stands for mereological inclusion):

> (1) G is *intrinsic* iff:
> for all $x < w < w'$, x has G in w iff x has G in w'.

Roundness is intrinsic, by this definition, because for any w including x and any elaboration w' of w, if it is round in either, it is round in the other. Accompaniment is extrinsic because w and its elaboration w' can overlap on x while differing in the amount of company they provide it.[19]

IX. ESSENTIAL EXTRINSICS

A consequence of (1) may give us pause. Call a property *absolutely essential* iff anything that has (lacks) it at all has (lacks) it in every world where it exists. Then *every absolutely essential property G is (1)-intrinsic.*[20] (If x has/lacks G essentially, there are not going to be *any* worlds in which x lacks/has G; so there are not going to be any that contain w as a part.) This result gives us pause because it falls afoul of well-known examples of essential properties due to Kripke and others.

Example 1: According to Kripke, I am essentially descended from a certain zygote Z; and it seems plausible as well that nothing can descend (in the relevant sense) from Z without being me. So *descending from Z* is an absolutely essential property. But on almost anybody's account, the zygote stopped existing before I started—whence *descending from Z* is an *extrinsic* property of mine. So how can (1) be believed when it tells us that extrinsicness and absolute essentiality are not compatible?

Example 2: Anything with the property of *being human* has that property essentially; and whatever lacks the property lacks it essentially as well. So *being human* is absolutely essential. But it is not intrinsic; how could it be when it involves an evolutionary lineage reaching back before the time not only of particular humans but of humanity itself? So *being human* is extrinsic and absolutely essential, again contrary to (1).

Example 3: Consider the property of *being me,* that is, *being SY.* Part of what it takes to be SY is to have the other two properties mentioned; so if they are extrinsic, *being SY* is too. But we know from the necessity of identity (and distinctness) that *being SY* is absolutely essential. So *being SY* is an extrinsic and absolutely essential property, in defiance of (1).

Where are we? Lewis, recall, had his problem of *accidental intrinsics*—the problem being that accidental intrinsic properties are, initial appearances to the contrary, not possible. He "solved" this problem by caving in to it—x's intrinsic properties do indeed have to be essential in one sense: x has them in every world where it exists—and then reestablishing his intended subject matter elsewhere—intrinsic properties do *not* have to be essential in the alternative sense that x shares them with all its counterparts.

Our problem is something like the opposite of Lewis's. It has to do not with accidental intrinsics but *(absolutely) essential extrinsics.* The problem is that (absolutely) essential extrinsic properties, despite making clear intuitive sense, are threatening to come out impossible. Our solution, in the next section but one, will have a similar shape to Lewis's. First, there'll be the caving in: if we limit ourselves to entities of such and such a type, then, yes, absolutely essential properties do have to be intrinsic. (But that is no paradox because entities of the type in question *should* have intrinsic essences.)

Second, there'll be the reconstitution elsewhere of our intended subject matter: if we let in entities of such and such *other* types, then absolutely essential extrinsics become once again possible.

X. ALL-NOT-OVERLAP IS BAD

I just gave *being SY* as an example of an essential property that was extrinsic to its bearer. I didn't say that *all* identity properties—all properties of the form *being x*—had to be extrinsic. Such a claim would not be plausible, and it has rarely been defended in philosophy.[21] If anything, the tendency has been to fall into the opposite error: the error of seeing *being x* as the intrinsic property par excellence. The right thing to say, with that error now exposed, is that some identity properties are intrinsic and others are not.

Take for instance a pair of protons, or two space-time points, or two identically shaped hunks of gunk. Do these have their (separate) identities intrinsically, or not? One can, of course, conceive of the first hunk being the hunk it is due, say, to its causal origins, which were different from the origins of the second hunk. But it takes imagination. The natural (and so presumably not absurd) thought is that the hunks are what they are as a purely intrinsic matter; each gets its identity from its gunk, and the gunk in hunk 1 is not numerically the same as the gunk in hunk 2. A venerable tradition even maintains that when you dig down to the ultimate building blocks of reality—the level of protons and the like, or perhaps a deeper level still—identity is *always* intrinsic. One may agree with the tradition or one may not. But intrinsic identity is certainly not an *incoherent* idea, and so we should take care to allow for it in our theory of intrinsicness.

Allowing for it will be difficult, if overlap is forbidden. The quickest way to see this is via a widely accepted principle of Lewis's connecting intrinsicness to duplication:

> (#) the intrinsic properties are the ones that never distinguish duplicates—including, crucially, duplicates in different worlds.

Suppose that overlap is prohibited. Then the property of *being x* distinguishes x from its otherworldly duplicates; if the duplicates were also x then we'd have overlap. It follows that whatever x may be, it has its identity extrinsically.

This is what I am calling problematic. A theory of intrinsicness should not predict right out of the starting gate that there is no such thing as intrinsic identity. But it will predict this if it accepts (#) and rejects overlap.

One response to the difficulty, proposed by Langton and Lewis, is to treat the theory as rendering no judgment in cases like this; the theory

defines intrinsicness for "pure" or "qualitative" properties, not "impure" properties like *being x*.

But we might have had hopes of a unitary treatment of intrinsicness that applied to pure and impure properties alike.[22] If so, then our choices are either to reject (#) or to reconcile ourselves to the possibility of overlap. There is nothing wrong with (#), so we need to make our peace with overlap.[23]

The goal is a unitary treatment of intrinsicness—a theory as comfortable with impure intrinsics like *being x* as it is with pure intrinsics. How are we to tell if our attempt at a unified theory has succeeded? It would help if we had examples of intrinsic-natured items to test it against. Let's adopt for purposes of this paper the view that worlds are made up of intrinsic-natured atoms *a, b, c, . . .*[24] And let's make it a condition of adequacy on our (eventual) theory that it recognize *being a* as intrinsic for each atom *a*.

XI. ESSENTIALISM AND MEREOLOGY

Back now to the issue we were dealing with before the digression. Analysis (1) may or may not succeed with accidental extrinsics; but to (absolutely) *essential* extrinsics, like *being human*, it is quite blind. All such properties are found by the analysis to be intrinsic.

Could it be that we have been pitching (1) to the wrong audience? A case can be made that the Kripke examples do not refute (1) so much as calling attention to one of its presuppositions. (1) is an analysis for *monolithic mereologists*—people who see no real competitors in the part/whole department to the relation of *mereological parthood.*[25] And a case can be made that the problem of essential extrinsics *does not arise* for the monotholic mereologist, because for her *there are no such properties.*

Why not, you ask? Among the laws of mereology is one called the *unique-sum principle:* take any objects $x_1, x_2, x_3, \ldots$ you like, there is a unique object $S(x_i)$—their *sum*—with all the x_is as parts and all of whose parts overlap the x_is. If we accept this principle, then it becomes plausible to hold that an absolutely essential property has got to be intrinsic. What is the essence of the x_is' unique sum going to be, after all, if not

(a) to exist in exactly the worlds containing each x_i
(b) to exist at exactly the space-time positions occupied by any x_i
(c) to have the x_is as parts, and
(d) whatever is necessitated by (a)–(c)?[26]

Kit Fine calls the sum defined by (a)–(d) the *aggregate* of the x_is, and remarks that "it has . . . often been supposed that aggregation is the only legitimate method of [summation]."[27] Whoever supposes that sums are aggregates and adopts our assumption above that every x is the sum, ultimately, of intrinsic-natured atoms, is well on the way to supposing that x's essential properties are one and all intrinsic to it. For they will find it hard to resist the following line of argument:

> Suppose that P is essential to x, where x is the sum of intrinsic-natured atoms x_i. Then P is necessitated by Q = the property of existing when and where the x_is do, with them as parts. Given the intrinsicness of *being* x_i, x possesses Q regardless of what goes on outside its boundaries. Hence it possesses P the same way; hence x has P intrinsically. Now, if a property essential to x is intrinsic to it, then a property that can *only* be had essentially (an absolutely essential property) can only be had intrinsically. A property that can only be had intrinsically, however, is an intrinsic property.[28] So every absolutely essential property is intrinsic.

The unique-sum principle thus sits very well with (1); for it gives us a basis on which to *reject* the absolutely essential, extrinsic, properties that were causing (1) so much trouble.

If you are with me this far, though, you will see that the unique-sum principle sits very *ill* with the Kripkean perspective. From that perspective it is going to seem unduly limiting to suppose that the essence of an x made up of $x_1, x_2, \ldots$ has got to be (a)–(d) above.

If we're going to have a sum $S(x_i)$ that exists in a world w iff *all* the x_is exist there, why not a sum $S^1(x_i)$ that exists iff *any* of the x_is exist? Why not a sum $S^2(x_i)$ whose existence in w requires only that *most* of the x_is exist? Why not a sum $S^3(x_i)$ which exists iff all and *only* the x_is exist?[29] Why for that matter should there not be *lots of* sums $S^\alpha(x_i)$, alike in "ordinary" respects but differing in which of their properties they have essentially? And why should not some of these sums have extrinsic properties in their essences—as indeed one of the sums just mentioned already does?

These questions have no good answers, I think, or at least none that the Kripkean can be expected to find convincing. This is why I say that the moral of the Kripke examples is not that (1) is wrong, but that Kripkeans are not monolithic mereologists and so not among its intended audience.

XII. INTRINSICNESS FOR THE REST OF US

Even if (1) can be defended as capturing *a* notion of intrinsicness—acceptable perhaps to monolithic mereologists—it remains to be seen how intrinsicness is to be defined for "the rest of us," in particular, those who are convinced by the Kripke examples that the properties essential to a thing need not be intrinsic to it, or intrinsic at all.

Suppose that x has an extrinsic property G essentially. Then purely

mereological changes in the part of w outside of x may well have the effect of removing x from the scene. So x will not be there to "witness" G's extrinsicness in w'. How then is G's extrinsicness to be brought out?

The answer has got to be that it is not x but a *related* object x' whose G-ness in w' concerns us. An example will show what I mean. Suppose for argument's sake that creationists have got it partly right. The world did indeed spring into existence in 4004 B.C., complete with fossils, archeological remains, and other traps for the unwary. But, and this is where creationists overreach themselves, there was no act of God involved; the transition from nothing to something was a complete and utter ex nihilo miracle.

All of that given, it seems to me (and for the sake of the example, please let it seem to you, too) that it is essential to the kind *rabbit* that it originated more or less spontaneously, and, in particular, that it did not evolve from earlier kinds. It follows that Floppy here, who is essentially a rabbit, is essentially of a non-evolved kind. *Being of a non-evolved kind* is extrinsic, though; so the fact that Floppy has it essentially means that the property of *being identical to Floppy* is extrinsic as well.

Now consider a world w' that prefixes to w a long and complicated evolutionary history; w' is the sort of world that Darwinian types (wrongly) *believe* themselves to inhabit. When we look at the candidates for Floppyhood in w', we find that every one of them sits at the terminus of a continuous and biologically plausible developmental path tracking back to earlier species. Since it is essential to Floppy *not* to sit at such a terminus, Floppy does not exist in w'. What is it, then, that bears witness in w' to the extrinsicness of *being Floppy*? It has to be the "Darwinized" creature Floppy′ that takes her place there.

Our challenge is to say how in general this x'—the object whose failure to exemplify G in w' reveals G as extrinsic—is to be identified. The outlines of the answer seem fairly clear: x' should be an object occupying exactly the piece or portion of w' that x occupied in w. But that still doesn't clue us in to the identity of the object x' that we want.

An advantage of the monolithically mereological approach was that it always made sense to speak of *the* occupant of the-x-portion-of-w-transplanted-into-w'. The disadvantage, of course, was that the occupant was x itself, leading to the unhappy (from a Kripkean vantage point) results just noted: results that led us to reconsider our attachment to monolithic mereologism.

But although mereology can be pushed too far, it stands to reason that we should try to hold on to as much of it as we can, compatibly with the Kripkean data we are trying to accommodate. Why not let the view be that worlds are mereological *at bottom*—at the level of their "stuff" or "matter"—with the non-mereological aspects superimposed? Then we can say that at the bottom level, where mereology reigns, essences are always intrinsic; while at the higher Kripkean levels, where any sort of property can be essential, mereology graciously steps aside. There may be various ways of implementing this divide-and-conquer strategy, but the following seven step plan looks attractive:

First, α-sums:

> The x_is have a number of sums $S^{\alpha}(x_i)$, exactly alike except in their essences, or (what comes to the same) their transworld careers.

Two, aggregates:

> One of these α-sums is the *aggregate* $S^{0}(x_i)$; its essence is to exist iff each of the x_is do, when and where any of the x_is do.

Three, pieces:

> x is a *piece* of y ($x <^{0} y$) iff there are things z_i such that y is the aggregate of all of them, and x is the aggregate of some of them.

Four, portions:

> The *stuff* of a world w is the aggregate of all its atoms. Pieces of w-stuff, that is, aggregates of some or all of w's atoms, are called *portions* of w.

Five, parts:

> x is a *part* of y ($x < y$) iff there are α and z_i such that y is the α-sum of all of the z_is, and x is the α-sum of some of them.[30]

Six, coincidence:

> x and y *coincide* iff they have exactly the same parts—equivalently, iff x is part of y and vice versa.

Seven, existence:

> To *exist* in a world w is to be (a) a portion of w, or (b) constituted by (strictly: coincident with) a portion of w.

As the seventh step suggests, coincidence is the notion we need, but the definitions read better when framed in terms of constitution. A word should thus be said about what it is for x to constitute y, and why it does little harm to substitute constitution for coincidence. The account I have in mind is a slight variant of one recently proposed by Judith Thomson.[31] The idea is that x, although possessed of the same parts as y, "hugs" those parts more closely than y does:

> x constitutes y iff

(a) x coincides with y,
(b) any part of x essential to it has parts that are not essential to y, and
(c) no part of y essential to it fails to have parts that are essential to x.[32]

Because pieces of world-stuff hug their parts so very closely—*all* their parts are (modulo the coincidence relation) essential to them—almost anything coinciding with one is also constituted by it.[33] This is why not much harm can come of writing "x constitutes y" instead of "x is coincident with y" in the definitions to follow. When push comes to shove, however, it is always coincidence we really have in mind.

XIII. INTRINSICNESS AND CONSTITUTION

All this time we have been grappling with a problem raised long ago for (1). The problem was that x may or may not carry over into w'. And if x does not exist in w', then it makes no sense to ask whether x is G in w iff x is G in w'. Our proposed solution was to say that the portion of w that constitutes x there—the x-portion of w—*does* carry over into w', where it constitutes an x' that *can* be assessed for G-ness in w'. If we modify (*) to take account of this solution, we get

(**) G is intrinsic iff:
for all $x < w < w'$, x has G in w iff x' has G in w'—where x' is whatever is constituted in w' by the x-portion of w.

An intrinsic property, in other words, is one that never distinguishes between x in w, on the one hand, and what x's constituting matter goes on to constitute in expansions of w, on the other. Since constitution is itself explained in terms of part/whole, this brings us *close* to the sought-after reduction of intrinsicness to part/whole and modality.

XIV. A FORK IN THE ROAD

But we are not there yet. The "whatever" in (**) hides a choice. There are going to be *lots* of things x' constituted in w' by the x-portion of w. Which of them have to be G for a verdict of intrinsicness?

The answers that come to mind are: *at least one* of them has to be G, and *each* of them has to be G. So there are two disambiguations of (**) to consider. These will be easier to read if we abbreviate "x' is constituted in w' by the x-portion of w" to "x' is a copy of x":

(2) G is intrinsic iff
for all $x < w < w'$, x has G in w iff *some* copy x' of x has G in w'.
(3) G is intrinsic iff
for all $x < w < w'$, x has G in w iff *every* copy x' of x has G in w'.

Before investigating how (2) and (3) differ, we should notice a way in which both of them improve on (1). They agree in allowing extrinsic properties to be essential.

Example: Suppose that w is the 4004 B.C. world described in section XII, imagined again as actual. And let w' be the more inclusive world that we wrongly *believe* to be actual. Our friend Floppy in w essentially possesses the property G of *not* sitting at the terminus of a long evolutionary history. But none of its copies in w' shares this property. It makes no difference, then, whether we go along with (2) in demanding that all of Floppy's copies be G, or with (3) in demanding that some of them be G. Either way, the demand is not going to be met. And so both definitions classify G as extrinsic.

XV. CATEGORICALS AND HYPOTHETICALS

Before trying to tease (2) and (3) apart, we need to fill in the background metaphysical picture some more, especially the all-important notion of coincidence. Coincidence theorists come in many shapes and sizes, but the relation I have in mind works as follows.

Imagine that we have before us the objects coincident in a world w with a portion p of w. These objects are, *categorically* speaking, just alike. To the extent that categorical properties fix perceptual appearance, they look just the same.[34] The differences between them are *modal*—one has essentially a property that another has only accidentally—or (to be a little more accurate) *hypothetical*—a matter of what they are like counterfactually, dispositionally, causally, or along any dimension that respects their behavior in other possible worlds.[35] By "casting our gaze" over these other worlds in imagination, we can "see" that the objects have different modal careers.

If this is how we understand coincidence, then (2) and (3) agree on the intrinsicness-status of categoricals. For let G be categorical. Then either *all* of the objects coincident in w' with x's w-stuff are G, or *none* of them is G. And so the hypothesis on the right-hand side of (2) is equivalent to the one on the right-hand side of (3). The question is, how do (2) and (3) compare in their classification of non-categoricals?

XVI. GENERIC VS. SPECIFIC INTRINSICNESS

Intrinsicness and categoricality are, of course, not the same. But there does seem to be a rather striking analogy between them. One can bring the analogy out like this: G is intrinsic iff whenever x has G in a world w, it has it in a manner insensitive to goings-on outside of x; and it is categorical iff whenever x has G in a world w, it has it in a manner insensitive to goings-on outside of w.

Someone who takes this analogy very seriously—someone who hears "outside of w" as standing to "outside of x" roughly as "outside of the galaxy" stands to "outside of this office"—will be tempted to conclude that sensitivity to circumstances outside of w is ipso facto sensitivity to circumstances outside of x. They will be tempted to conclude, in other words, that *hypothetical properties can never be intrinsic.* Call the notion of intrinsicness on which this is so—on which intrinsics have to be categorical—the *generic* notion. ("Generic" because it lumps other worlds in with other places as sources of extrinsicness.)

Whether they arrive at it by the indicated route or not, it is the generic notion that writers on intrinsicness often have in mind. Peter Vallentyne, for instance, says that "water-solubility and the like might seem like intrinsic properties, but once one recognizes [their] dependence on what the laws of nature are, it seems more correct to classify such properties as extrinsic."[36] After all, Vallentyne argues, the laws of nature "are part of 'the rest of the world'"—they are sensitive to changes in the world outside the object. If a lump of sugar can lose its solubility just by being placed in different surroundings, then solubility is not intrinsic.

Even if we agree, as not everyone does, that the laws in force at a given location depend on goings-on in "the rest of the world," Vallentyne's argument can be resisted. Dispositions are dispositions to behave in particular ways *in particular circumstances.* But if laws are circumstantial in the way that Vallentyne thinks, then there is no obvious reason why the circumstances should not be understood to include them.

An extrinsic disposition D_1 can almost always be made into a (more) intrinsic one D_2 by taking the external factors which make D_1 come and go, and loading them into D_2's triggering circumstances. This is one way—a very old-fashioned way, to be sure—of understanding the difference between weight and gravitational mass. An object's weight, on the old-fashioned account, is its disposition to depress a properly constructed scale so as to elicit a reading of so many pounds *in the local gravitational field, weak or strong as that field might be.* This nicely explains why weight is extrinsic: the reason your weight changes when you go to the moon, even though intrinsically you remain just the same, is that a different field becomes local. An object's mass, on the other hand—its disposition to elicit a reading of so many pounds *in a gravitational field of such and such strength, wherever a field of that strength might be*—is not extrinsic, or anyway not *as* extrinsic as its weight.

All of this is to suggest that the generic notion of intrinsicness is not the only one we understand, or the only one we have need of.[37] Sydney Shoemaker in a celebrated paper distinguishes two types of causal power:

> A particular key on my key chain has the power of opening locks of a certain design. It also has the power of opening my front door. It could lose the former power only by undergoing what we would regard as a real change, for example, a change in its shape. But it could lose the latter without undergoing such a change; it could do so in virtue of the lock on my door being replaced by one of a different design. Let us say that the former is an intrinsic power and the latter is a mere-Cambridge power.[38]

"Mere-Cambridge" powers are in later writings called "extrinsic" powers.[39] The point either way is the same: if the key's power to open a lock of such and such a type is intrinsic to it, then clearly we need a version of intrinsicness on which it does not entail categoricality. Since the reverse entailment does not hold in anyone's book, we want intrinsic/extrinsic to crosscut categorical/hypothetical in every possible way.

Another example to consider here is identity. Atoms have their identities as an intrinsic matter. But *being* x is always hypothetical, since one can imagine a distinct thing x^* that is indiscernible from x when we bracket their counterfactual careers (x^* might differ from x just in the fact that it exists in fewer worlds).[40] The property of *being a,* where a is an atom, is thus intrinsic and hypothetical.

Other examples could be mentioned, but they aren't really needed: it is obvious that we have, in addition to a generic notion of intrinsicness that entails categoricality, a *specific* notion that leaves an intrinsic property's status as categorical undecided. The specific notion is the more discriminating and hence (it seems to me) the more useful one. A generically intrinsic property is just a specifically intrinsic one that happens also to be categorical; defining specific intrinsicness in terms of generic would not be so easy.

XVII. INTRINSICNESS$_3$ = GENERIC INTRINSICNESS

Suppose that the space of particulars is *full:* for any assignment F of (coinstantiated) categorical properties to worlds, there is something x_F existing in just the worlds in F's domain and exemplifying in each of those worlds the properties that F assigns to it. Then, since a property is hypothetical iff it can be made to vary by manipulating otherworldly categorical profiles, (3) makes all hypothetical properties extrinsic.

Why does (3) have this result? Setting w' in (3) equal to w, we see that G is intrinsic in the sense of (3) only if it never distinguishes coincidents in the same world. No hypothetical property can meet this condition. If H is hypothetical, then there is a world w containing an x such that (i) x is H in w, but (ii) an x' differing from x only in its categorical properties in other worlds would not be H in w. By fullness, such an x' is bound to exist. So for any hypothetical property H, H's hypotheticality is "witnessed" by a pair of H-discernible coincidents in some world w. It follows that H is not intrinsic.

Example: Imagine a statue Goliath composed of a hunk of wax Lumpl. Goliath is essentially of a certain intrinsic shape—it is essentially, let us say, so-shaped. But to be so-shaped is only accidental to Lumpl, which might never have been formed into a statue at all. Since Goliath is coincident with Lumpl, it follows that Goliath in w is a copy of something (Lumpl) in w that differs from Goliath in point of essential so-shapedness. So the property of being essentially so-shaped is not intrinsic.

XVIII. INTRINSICNESS$_2$ = SPECIFIC INTRINSICNESS

Are hypothetical properties extrinsic in the sense given by (2)? Sometimes yes, sometimes no. Consider again the property of being essentially so-shaped. The reason this came out extrinsic$_3$ was that not everything coincident with Goliath in w' was essentially so-shaped; Lumpl, for one, was so-shaped only accidentally.

But all that (2) asks is that *something* coincident with Goliath be essentially so-shaped. And given fullness, this condition is met for any world w' containing p = the Goliath-portion of w. For consider a super-Goliath stipulated to possess exactly the categorical properties of p in w' and to exist in no other worlds. Super-Goliath is an essentially so-shaped coincident of p. And that's precisely what is needed for essential so-shapedness to come out intrinsic$_2$.

XIX. DOING IT WITH DUPLICATES

A property is intrinsic, according to Lewis, iff it never distinguishes duplicates. This reduces the problem of defining intrinsicness to the problem of explaining duplication. Suppose that we like this reduction, but don't like Lewis's explanation of what it is for x and y to be duplicates.[41] If we could obtain duplication from part/whole, that would be a move in the right direction.

Above we defined x' in w' to be a *copy* of x in w iff (a) w is part of w' (this to ensure that the x-portion of w persists into w'), and (b) x' is constituted in w' by the transplanted x-portion. The copying relation joined with its converse will be called *immediate duplication. Duplication* is the ancestral of immediate duplication: the relation that x_1 in w_1 bears to x_n in w_n iff there is a string of x_is and w_is ($1 \leq i < n$) such that x_{i+1} in w_{i+1} is an immediate duplicate of x_i in w_i.

Now we are almost there; all that remains is to distinguish two different kinds of transworld similarity to serve the needs of our two definitions. Say that x in u and y in v *agree* on a property G iff: x has ±G in u iff y has ±G in v. And say that they *concur* on G iff: x has coincidents with ±G in u iff y has coincidents with ±G in v. Then the following are equivalent to (2) and (3):

(2′) G is intrinsic iff duplicates always concur on G.
(3′) G is intrinsic iff duplicates always agree on G.

Exactly as before, (2′) gives us a modally neutral notion of intrinsicness—what above we called the specific notion—while (3′) defines the generic notion whereby an intrinsic property has got to be "modally intrinsic" or categorical.

XX. CONCLUSION

Lewis in "Extrinsic Properties" locates intrinsicness in "a tight little family of interdefinables," and muses about the possibility of breaking out. Later, in "New Work for a Theory of Universals," he sees his opening: reduce intrinsicness to duplication, and (the break-out point) duplication to naturalness. Still later, in "Defining 'Intrinsic'," he and Langton decide that they can make do with just an aspect or component of naturalness, viz., (relative) nondisjunctiveness.

But while that is certainly progress, neither of the proposed reducers—naturalness, nondisjunctiveness—*feels* like it is has much to do with "what intrinsicness is." Another way to put it is that both of the Lewis definitions trade on de facto connections with intrinsicness rather than de jure ones. It might be doubted, of course, whether intrinsicness *has* any interesting de jure connections with other notions. But conditionals like

(%) x is a part of y only if x cannot change intrinsically without y changing intrinsically as well.

show that this pessimism is unjustified. The de jure appearance of (%) led us to attempt a broadly mereological account of intrinsicness. Whether the account gives the right results in all cases, it does seem to lay the issue at a doorstep more in the right neighborhood.

The conditional (%), although putting us on the track of an account less accidental-feeling than Lewis's, also suggests a way in which our approach might be thought to fall short of his. Naturalness has no chance whatever of being explained in terms of intrinsicness; so a successful analysis of intrinsicness in terms of naturalness would probably amount to a reduction. One would truly have broken out of Lewis's "family of interdefinables."

An analysis in terms of part/whole, however, might or might not constitute a reduction, depending on whether part/whole was definable reversewise from intrinsicness. One is certainly accustomed to thinking of part/whole as the more basic of the two, but a look at (%) makes me wonder. Is a reversewise definition possible? If so, then we have at best enlarged Lewis's circle a little. If, on the other hand, part/whole deserves its reputation as primitive and undefinable, then perhaps we (too) can claim to have broken out.

APPENDIX

This appendix explores how far we can disentangle ourselves from the modal realism presupposed in the main body of the paper. I'll suppose that we are confronted with a series of three objectors. ARG says that the *argument* we gave for our account of intrinsicness presupposed modal realism. PAL says that argument or no argument, the account is not *plausible* without modal realism. EXP maintains that our account is not even *expressible* without modal realism.

ARG

A lot is riding on the idea that if one world is part of another, then this makes for an area of intrinsic match-up between them. The argument for this turned on Lewis's problem of accidental intrinsics. But as Lewis would be the first to admit, accidental intrinsics are a problem only for modal realists—only for those conceiving counterfactual worlds as big concrete objects on a par with the actual world. Modal realism is not a hugely popular doctrine.[42] How do you propose to defend your idea to the rest of us?

Reply: This misrepresents the argumentative strategy. Accidental intrinsics came in not to *support* any claim about parts and intrinsicness, but to quiet a certain worry that the modal realist might feel. I said: look, to the extent that your opposition to overlap is based on the problem of accidental intrinsicness, you needn't be worried about overlap of the specific sort envisaged here. That one world can be part of another wasn't argued for at all; it was taken for granted. Also taken for granted was the claim that if one thing overlaps another, there is going to be complete intrinsic similarity with respect to the shared part. Nobody would question this in the intraworld case; why would the transworld case be different?

PAL

I'll tell you why (it would be different in the transworld case). Not all transworld relations are created equal. On the one hand, you've got "genuine" transworld relations like *being the same color as*. These really do inherit the features of their intraworld originals. Transworld *same color as*, is transitive and entails *being colored*, just like the intraworld version.

On the other hand, you've got "degenerate" transworld relations like *touching*. All it can mean to say that x in w touches y in w' is that when you get them together in the same world, say v, you find that x touches y there. Degenerate transworld relations do *not* inherit the features of their intraworld originals; transworld touching, for example, is compatible with toucher and touchee each being quite alone in their original worlds.

Now if modal realism is correct, then just *maybe* there is a non-degenerate relation of transworld parthood. Otherwise though, parthood is like touching; when you try to apply it between worlds, it goes degenerate.[43] If that is right, then the fact that intraworld parts have a region of intrinsic match with their containing wholes says nothing whatever about the transworld case.

Reply: Transworld parthood is *not* degenerate. If it were, then (i) to be transworld part of y, x would have to be intraworld part of y (in some salient world), and (ii) if x were intraworld part of y (in some salient world), x couldn't *avoid* being transworld part of y. I will argue that neither (i) nor (ii) is at all plausible.

Against (ii): The Kiwanis Picnic would have been considerably shorter had the world popped (miraculously!) out of existence just as the soda was being opened. Indeed, more is true: the picnic that would have been is a *proper part* of the picnic that actually was. According to (ii), though, this is impossible, for the picnic that was is a part of the picnic that would have been.

Why is that? Well, the picnic that would have been is none other than the picnic that actually took place; it is that picnic as it might have been rather than as it is. But if they are the same, then there is no question but that each intraworld includes the other; a thing always intraworld includes itself. By (ii), finally, this intraworld inclusion suffices for transworld inclusion. *The denier of genuine transworld parthood can thus make no sense of the idea that the picnic that would have been—the actual picnic with its tail end chopped off—is a proper part of the actual picnic.*

Against (i): The Rainbow Rally includes what *would* have been the

entire Spartacus League demonstration—would have been, if a few Spartacist rowdies wrongly supposed to be out of town had not turned up to demonstrate *against* the Rally. According to (i), the *would-be* demonstration is part of the Rally only if the *actual* demonstration is. So, although the would-be demonstration precisely *omits* the rowdies, it still has to pay the mereological price for their behavior. That seems absurd; *the Spartacist demonstration as it would have been is wholly included in the Rally.*

Degenerate transworld relations are a dime a dozen. Any "ordinary" relation R gives rise to a bunch of them by the formula: x in w bears transworld R^v to y in w' iff x bears R to y in v. But these degenerate R^vs are almost never what is meant by "transworld R." It would be one thing if there were no alternative to the degenerate interpretation. But in this case there clearly is: it's the interpretation we come to naturally when, e.g., we dispute with Lewis about whether there are counterfactual worlds with the actual world as a part. If "part" here stood for the degenerate relation, the answer would be obvious; it isn't, so it doesn't.

EXP

Your account is committed to counterfactual worlds by virtue of explicitly quantifying over them. And the worlds you quantify over had better be Lewisian concrete worlds. Because it seems very doubtful that ersatz worlds, e.g., sets of propositions, are going to stand in the requisite inclusion-relations.

Reply: You're right that the worlds of (2) and (3) are concrete. But just maybe the analysis can be (re)construed as speaking of worlds that, although they *could* have existed, as it happens do *not* exist—not even in a place bearing no spatiotemporal relations to the speaker.

The reason that the analysis appears to require existent concrete worlds is that it makes essential play with transworld relations such as transworld part/whole. How can the counterfactual picnic be (transworld) part of the actual one unless both are somehow *there,* in their containing worlds?

This is a fair question, but it has an answer; it had *better* have one, for colloquial English is thick with talk of transworld relations, and other worlds do not seem to come into the picture. You might say, for example, that

> (i) the car I would have had, if I'd installed afterburners in my old 'Vette, is faster than the Camaro I do have.

Does this commit you to the existence of a (counterfactual world containing a) counterfactual car, fitted out with afterburners as your actual car is not? Not at all. Your claim in essence is that

> (i′) it could have happened that I had a 'Vette that was faster than my Camaro is in actual fact.

And (i′) makes no mention of counterfactual objects except in the scope of a modal operator. What remains to be seen is whether the same can be done with (2) and (3), that is, whether they too can be restated in a way that avoids wide-scope quantification over things that (modal realism aside) do not exist.

Before attempting this, let me concede right off that a formulation as colloquial as (i′) is going to be hard to come up with. Our everyday modal devices are quickly pushed to their limits when the transworld comparisons become too involved, as in

> (ii) the car I would have had if I'd installed afterburners in the 'Vette is intermediate in speed between the one I would have had if I'd supercharged my Camaro and the Camaro as it is actually.

About the best we can do with this is

> (ii′) it could have happened that after supercharging my Camaro was such that [it could have happened that I had a 'Vette with afterburners such that {the Camaro was faster than the 'Vette was and the 'Vette was faster than the Camaro *is*}]

This may not be wonderful English, but it seems close enough in *spirit* to (i′) that it would be strange to discern here a qualitative change in subject matter—so that while (i′) gave a partly modal description of actuality, (ii′) described non-actual, merely possible, items, viz., other possible worlds and their inhabitants.

If that is right, then our next step should be to seek a Loglish-type regimentation of the language used in (i′) and (ii′), in the hope that it permits a noncommittal reformulation of transworld-part talk, and ultimately a noncommital reformulation of (2) and (3). A number of people have applied themselves to this sort of problem,[44] and they have had considerable success with a device called "multiple indexing." Rather than trying to explain the idea from scratch, let me illustrate it with translations of our two target sentences. Think of "c" as a proper name of my actual Camaro, and "B" and "S" as standing for the relevant sort of afterburner-enhanced 'Vette and supercharged Camaro:

(i′′) possibly $(\exists y)$ (By & actually^c [y is faster than c]).
(ii′′) possibly$_1$ $(\exists z)$ ($z = c$ & Sz & possibly$_2$ $(\exists y)$(By & actually$_1^z$ actually$_2^y$ actually^c [z is faster than y is faster than c]))

Now let's try the method out on

> (iii) yesterday's Kiwanis picnic was only a part of the picnic we would have had if not for the interruption.

Here it is, nearly enough, in noncommittal English:

(iii´) it could have happened that we had a picnic of which the picnic we did have was a proper part.

And here it is in Loglish, with *"k"* standing for the picnic and "P" describing the relevant sort of prolonged picnic:

(iii´´) possibly $(\exists y)$ (Py & actually^{k} [k is a proper part of y]).

Now let's see what can be done with (2) and (3). I'll assume that, necessarily, there is one and only one world (which is not, of course, to say that "the world" is rigid). The letters *"u"* and *"v"* are world-variables, *"p"* and *"q"* range over world-portions, i.e., aggregates of atoms, "≈" expresses the coincidence relation, and "<" stands for parthood:

(3´´) G is *intrinsic* iff:
necessarily$_1$ $(\forall u)$ $(\forall p < u)$ $(\forall x \approx p)$ necessarily$_2$ $(\forall v)$ $(\forall q < v)$ $(\forall y \approx q)$ (actually$_1^{pu}$ actually$_2^{qv}$ $(p = q < u < v) \rightarrow$ [actually$_1 x$ G$x \leftrightarrow$ actually$_2^{y}$ Gy])

G is intrinsic$_3$, in other words, when the following holds necessarily: if a mass of atoms that composes x would, had a more inclusive world obtained, have composed y, then x is G iff y would have been G. The definition of intrinsicness$_2$ is nearly the same, except that the two coincidence-quantifiers get moved to either side of the final biconditional:

(2´´) G is *intrinsic* iff:
necessarily$_1$ $(\forall u)$ $(\forall p < u)$ necessarily$_2$ $(\forall v)$ $(\forall q < v)$ (actually$_1^{pu}$ actually$_2^{qv}$ $(p = q < u < v) \rightarrow$ [actually$_1$ $(\forall x \approx p)$ $\pm$G$x \leftrightarrow$ actually$_2$ $(\forall y \approx q)$ $\pm$Gy])

The ±Gs towards the end are to indicate a conjunction of two biconditionals, one with G unnegated on both sides, one with it negated. Translated into English, what (2´´) says is that G is intrinsic$_2$ iff necessarily, any mass of atoms that composes only Gs (non-Gs) would still have composed only Gs (non-Gs), had a world obtained of which the actual world is only a part.

NOTES

1. This paper repays an old debt to Sydney Shoemaker. Years ago, as editor of *Philosophical Review*, he allowed me to cite a manuscript on intrinsicness that I was unable to provide at the time. He asked for it again in 1994, when he taught a seminar on intrinsicness; all I could offer was a one-page sketch of the main idea. Now you've got it, Sydney! I had help from Sally Haslanger, Ned Hall, Rae Langton, Jennifer McKitrick, Gideon Rosen, Alan Sidelle, Ted Sider, Judy Thomson, and Ralph Wedgwood. I wish I could acknowledge a greater debt to Peter Vallentyne's "Intrinsic Properties Defined"; it would have saved me a lot of trouble to have seen it earlier. An implicit subtheme here is that there is more to Vallentyne's approach than Langton and Lewis give him credit for.

2. Jaegwon Kim, "Psychophysical Supervenience," *Philosophical Studies* 41 (1982): 51–70.
3. Ibid.
4. David Lewis, "Extrinsic Properties," *Philosophical Studies* 44 (1983): 197–200.
5. "New Work for a Theory of Universals," *Australasian Journal of Philosophy* 61 (1983): 343–77. The formulation in the text leaves out some irrelevant complications.
6. "Accident" is meant to be compatible with a fact's holding necessarily. If God and the null set are both necessary beings, then it is impossible for either to exist without the other. Still the correlation is accidental in the sense I have in mind.
7. David Lewis and Rae Langton, "Defining 'Intrinsic'," *Philosophy and Phenomenological Research* 58 (1998).
8. David Lewis, *On the Plurality of Worlds* (London: Blackwell, 1986), 87–88.
9. Ibid., 88.
10. Ibid., 199.
11. Ibid., 199–200. This is not the first occurrence of "How many fingers?" in world literature; O'Brien puts the question to Winston in *1984*.
12. Meaning the ones he has in every world that contains him.
13. Some might say that worlds cannot vary extrinsically either. But our resistance to extrinsic variation is much weaker, or the branching conception of worlds would not be so popular. See John Bigelow, "The World Essence," *Dialogue* 29 (1990): 205–17, for the view that a world's *intrinsic* nature is largely accidental to it.
14. Advocates of the branching conception appear to take this for granted. A counterfactual world w' branches from the actual world @ at time t iff world w, which has the history of @ up to time t as its entire history, is an initial segment of w'. That the relevant intrinsic features of @ "carry over" into w and thence into w' goes without saying.
15. To go by a recent report in *Scientific American*, this is not altogether out of the question. Atom-smashers have felt moved on occasion to calculate the chances of their experiment flipping the universe into "null-state."
16. Lewis has other objections to overlap that are meant to apply even when there is no variation in intrinsic properties. He says, for instance, that overlap in the form of branching "conflicts with our ordinary presupposition that we have a single future" (*Plurality of Worlds*, 207). But these other objections are not advanced as decisive, and if overlap opens the door to a new account of intrinsicness, that would be an argument on the other side.
17. At least, none of relevance to us; set-membership is a potential counterexample that I would rather put aside.
18. Actually, perhaps it is not so clear. If x sits at a boundary of w, it could happen that x has a different shape in the larger w' than it had in x. Concessions about to be made in response to a different worry should alleviate this one as well.
19. Does (1) make a material object's location intrinsic to it? I don't think so. The problem arises only if one is a substantivalist about space-time; it's only if location is the property of occupying such and such space-time points that x's location can be expected to "follow it" from w into the more inclusive w'. But for G to be intrinsic, it's not enough that it can't be lost through world-expansion; a further requirement is that it can't be gained that way either. And this gives us a way out. No one supposes that the space-time points a table occupies are to be counted among the table's *parts*. So the table should be able to survive intact into a world v from which those points had been removed. Consider now v and the more inclusive world v' with which we began. That the table *gains* the property of occupying such and such space-time points in the transition from v to v' tells us that location is not intrinsic according to the definition in the text. (See also n. 30 below.)
20. In this paper, "x has G essentially" means that x has G in every world where it exists. The door is left open to a counterpart-theoretic rejoinder to the argument of the main text.

21. But see Graeme Forbes, *Metaphysics of Modality* (Oxford: Clarendon Press, 1985), and my review in *Journal of Philosophy* 85 (1988): 329–37.
22. Lloyd Humberstone calls this intrinsicness-as-interiority. See his "Intrinsic/Extrinsic," *Synthese* 108 (1996): 205–67.
23. This is made easier by the fact that the principal objection to it doesn't apply to the kind of overlap we are contemplating.
24. Not, I hasten to add, the same ones in every world; some worlds share atoms, others do not. Being composed of atoms $a_1 \ldots a_n$ will come out intrinsic on our account. See Humberstone, "Intrinsic/Extrinsic," on the distinction between intrinsicness and nonrelationality.
25. So-called because it is defined in large part by the fact that it satisfies the laws of mereology.
26. Kit Fine, "Compounds and Aggregates," *Nous* 28 (1994): 137–58. Compare also Judy Thomson's "some-fusions," in Thomson, "The Statue and the Clay," *Nous* 32 (1998): 149–73.
27. Fine, "Compounds and Aggregates," 138.
28. The relation between intrinsic *properties* (our topic here) and intrinsic *predication* (having a property intrinsically) is nicely elaborated by Humberstone.
29. That is, a sum $S^3(x_i)$ that exists iff all the x_is exist, and nothing disjoint from them exists.
30. Special cases aside, x is a part of y iff it occupies a subset of the space-time positions occupied by y. The special cases include, for instance, a bundle of compresent tropes vis à vis some particular trope t in the bundle; the bundle may occupy the same space-time position(s) as t, but it isn't part of t. Another example might be Casper the ghost passing through a mountain or a larger ghost. The special cases matter. An account of intrinsicness ought to distinguish the trope's intrinsic properties from those of the co-located bundle; likewise Casper and the co-located mountain-part. Intrinsicness is not a spatiotemporal notion except per accidens—the accident being that part/whole tends to be spatiotemporal.
31. Thomson, "The Statue and the Clay."
32. At the risk of oversimplifying, the difference between this definition and Thomson's is that she has "some" at the beginning of (b) where I have "any." This has the result, which I find unwelcome and she does not, that if Lump1 constitutes Goliath, then take any z you like (e.g., the planet Saturn), the fusion of Lump1 with z constitutes the fusion of Goliath with z.
33. For an example of something coincident with a piece x of world-stuff in w that x does *not* constitute in w, consider a y that is "just like" x in w but exists in no other worlds. *Everything* about y is essential to it, so condition (b) cannot be met.
34. Note that the compound of the x_is may well be categorically different from their aggregate. It is likely to have a shorter life since it goes out of existence when any of the x_is ceases to exist. This shows that Fine's compounding is not a sum operation in my sense (likewise Thomson's operation of all-fusion).
35. I am not using "categorical" and "hypothetical" to help define "intrinsic." I am using them to work out the consequences of definitions (2) and (3) for people who, like me, think that coincidents (things with the same parts) are categorically alike and distinguished by their hypothetical properties.
36. Peter Vallentyne, "Intrinsic Properties Defined," *Philosophical Studies* 88 (1997): 209–19.
37. Some will see a different problem with the generic notion. They will object that the generic notion, far from being the *only* legitimate one, is in fact *illegitimate*, because dispositional properties are always intrinsic. See, e.g., David Lewis in "Finkish Dispositions," *Philosophical Quarterly* 47 (1991): 142–58 and George Molnar, "Are Dispositions Reducible?" *Philosophical Quarterly* 49 (1999): 1–17. Both the always-extrinsic and the always-intrinsic positions are rejected here as too extreme.

38. Shoemaker, *Identity, Cause, and Mind* (London: Cambridge University Press, 1984), 221.
39. E.g., at 105–6 of *The First-Person Perspective and Other Essays* (London: Cambridge University Press, 1996).
40. See my "Identity, Essence, and Indiscernibility," *Journal of Philosophy* 84 (1987): 293–314.
41. Lewis relies on a distinction between natural and unnatural properties that he treats as primitive. One can agree that a primitive notion of naturalness is needed in philosophy, without agreeing that it is needed for the definition of intrinsicness, or that its availability is enough of a reason to use it.
42. Not to mention that it has surprising consequences for intrinsicness. Assuming modal realism, accompaniment (as opposed to accompaniment by something you bear spatiotemporal relations to) comes out intrinsic; duplicates never differ with respect to it because *everything* is accompanied. Accompaniment has traditionally been the paradigm of an extrinsic property.
43. For discussion, see Nathan Salmon's *Reference and Essence* (Princeton: Princeton University Press, 1981), 121ff.
44. See Forbes, *Metaphysics of Modality*, 89ff. and references there.