

David Lewis on Persistence¹

Katherine Hawley University of St Andrews

Next, persistence through time. I take the view that nothing endures identically through time. (Except universals, if such there be; their loci would coincide with relations of qualitative match, would indeed constitute these relations, so they would commit no violations of Humean Supervenience.) Persisting particulars consist of temporal parts, united by various kinds of continuity. To the extent that the continuity is spatiotemporal and qualitative, of course it supervenes upon the arrangement of qualities. But the continuity that often matters most is causal continuity: the thing stays more or less the same because of the way its later temporal parts depend causally for their existence and character on the ones just before. So the spatiotemporal boundaries of persisting things, for example people, can supervene on the arrangement of qualities, provided that causation does. (*Philosophical Papers, volume II, xiii*)

To persist is to exist at more than one time, to transcend the momentary. How do things achieve this? We might answer with talk of thermodynamic stability, molecular bonds, photosynthesis, the porcupine's spines, German manufacturing standards, legal protection of ancient monuments, or the uncanny ability of children to extract care from their parents. In Lewis's terms, such answers explain the existence of spatiotemporal and qualitative continuities over time in causal terms, by reference either to the causal mechanisms which directly underpin such continuities, or to their preconditions and external circumstances. Explanations may differ according to the kind of object in question: German washing machines and yew trees are both long-lasting, relative to other types of appliance or tree respectively, but the reasons for their longevity are quite different.

The metaphysicians have a further question about persistence, a question which is expected to have the same answer for all sorts of concrete objects. What is it for

¹ Many thanks to Jonathan Schaffer for his comments on an earlier draft. This work was supported by the European Community's Seventh Framework Programme FP7/2007-2013, under grant agreement no. FP7-238128.

something to exist at more than one time? For Lewis, a thing exists at more than one time by having distinct stages, each of which is located at a different time. These stages are parts of the persisting object: added together, they *are* the persisting object. On this view, persistence through time is analogous to extension through space: a spatially-extended object occupies more than one point at a single time by having different spatial parts located at different places. Likewise, for Lewis, a temporally-extended object occupies more than one time by having different temporal parts located at different times. “Persisting particulars consist of temporal parts, united by various kinds of continuity.”

The intended contrast is the view that concrete things “endure identically through time”, as universals do if they exist. Universals are wholly located where their instances are: the universal *having mass 1kg* is entirely present in each of the 1kg bags of flour on the supermarket shelf, and that very same universal will be present in future bags of flour, just as it was in past bags of flour. The universal does not portion itself out, a part here and now, a part there and then: instead, the whole universal is in each place it is needed. In the same way, an enduring concrete thing, if such there be, would not portion itself out over time, a stage then and a stage now: instead, the whole persisting object is located at each time of its existence.

Following Lewis (who credits Mark Johnston), these rival views of persistence are now known as ‘perdurantism’ and ‘endurantism’ respectively:

Let us say that something *persists* iff, somehow or other, it exists at various times; this is the neutral word. Something *perdures* iff it persists by having different temporal parts, or stages, at different times, though no one part of it is wholly present at more than one time; whereas it *endures* iff it persists by being wholly present at more than one time. Perdurantism corresponds to the way a road persists through space; part of it is here and part of it is there, and no part of it is wholly present at two different places. Endurantism corresponds to the way a universal, if there are such things, would be wholly present wherever and whenever it is instantiated. (*On the Plurality of Worlds* 202).

(Lewis had previously used ‘endure’ as the neutral word (‘Paradoxes of Time Travel’, p.68 in *Philosophical Papers volume II*)).

Lewis touched upon issues of persistence throughout his publishing career, from ‘Counterparts of Persons and Their Bodies’ in 1971, to ‘How Many Lives has Schrödinger’s Cat?’ in 2004. The most extensive discussions can be found in ‘The Paradoxes of Time Travel’, ‘Survival and Identity’ and its later postscripts, *On the Plurality of Worlds*, ‘Rearrangement of Particles: Reply to Lowe’, and ‘Tensing the Copula’; he remained committed to perdurance theory at every stage. In this chapter, I will explore the connections between Lewis’s perdurance theory and his Humean Supervenience, arguing that his influential argument about temporary intrinsics is best seen in this light. I then turn to a domestic dispute within the anti-endurantist camp: why does Lewis identify ordinary objects with world-bound parts of transworld objects, but not with time-bound parts of transtemporal objects? Given that Lewis is a counterpart theorist about modality, why isn’t he a stage theorist about persistence?

Section 1: Persistence and Humean Supervenience

Humean Supervenience is yet another speculative addition to the thesis that truth supervenes on being. It says that in a world like ours, the fundamental relations are exactly the spatiotemporal relations: distance relations both spacelike and timelike, and perhaps also occupancy relations between point-sized things and spacetime points. And it says that in a world like ours, the fundamental properties are local qualities: perfectly natural intrinsic properties of points, or of point-sized occupants of points. Therefore it says that all else supervenes upon the spatiotemporal arrangement of local qualities throughout all of history, past, present and future. (‘Humean Supervenience Debugged’, 225-6 in *Papers on Metaphysics and Epistemology*.)

When philosophers claim that one or another commonplace feature of the world cannot supervene on the arrangement of qualities, I make it my business to resist. (*Philosophical Papers volume II*, xi).

The points of resistance are laws of nature, counterfactuals, causation, mind and language, chance, and persistence. Lewis twice discusses possible worlds exemplifying certain persistence patterns, using them as examples of distant worlds which violate Humean Supervenience (he is committed to Humean Supervenience only for the actual world and its neighbours):

Two worlds might indeed differ only in unHumean ways, if one or both of them is a world where Humean supervenience fails. Perhaps there might be extra, irreducible external relations, besides the spatiotemporal ones; there might be emergent natural properties of more-than-point-sized things; *there might be things that endure identically through time or space, and trace out loci that cut across all lines of qualitative continuity*. It is not, alas, unintelligible that there might be suchlike rubbish. Some worlds have it. And when they do, it can make differences between worlds even if they match perfectly in their arrangements of qualities. (*Philosophical Papers volume II*, x, my italics.)

...one way to get a difference between worlds with the exact same arrangement of local qualities is to have things that are bilocated in spacetime. Take two worlds containing spheres of homogeneous matter, unlike the particulate matter of our world; in one world, the sphere spins and in the other it doesn't; but the arrangement of local qualities is just the same. These are worlds in which things persist through time not by consisting of temporal parts, but rather by bilocation in spacetime: persisting things are present in their entirety at different times. The difference between the spinning and the stationary spheres is a difference in the pattern of bilocation. ('Humean Supervenience Debugged', 227 in *Papers in Metaphysics and Epistemology*.)

The thought seems to be this: we can maintain that Humean Supervenience is true at the actual world and its neighbours only if we accept that perdurance theory is true at the actual world and its neighbours. Given Lewis's determination to resist philosophical challenges to Humean Supervenience, this counts in favour of perdurance theory as a description of the actual world. But how exactly would enduring objects violate Humean Supervenience? Segments of the homogeneous spheres, which "trace out loci that cut across all lines of qualitative continuity", clearly violate the claim that all facts supervene upon those facts about qualitative continuity. But endurance theorists needn't be committed to the possibility of such cross-cutting objects. Here are two versions of endurance theory which seem compatible with Humean Supervenience:

Plenitudinous endurance theory: all and only timelike sequences of matter-filled points are trajectories of enduring objects.

Qualitatively-restricted endurance theory: all and only those maximal timelike sequences of matter-filled points which exhibit appropriate qualitative and causal continuity are trajectories of enduring objects.

Each of these theories seem to be compatible with Humean Supervenience. That is, each theory is such that, if the actual world and its neighbours satisfied that theory, then Humean Supervenience would be true of those worlds. Any two worlds which match in their arrangements of local qualities and causal relations will also match in their distribution of (i) timelike sequences of matter-filled points and (ii) maximal timelike sequences of matter-filled points which exhibit appropriate qualitative and causal continuity. Given plenitudinous endurance theory, there is no opportunity for patterns of bilocation to become detached from lines of continuity, since patterns of bilocation are ubiquitous. Given qualitatively-restricted endurance theory, enduring objects are not free to cut across lines of qualitative-causal continuity: they are somewhat like the multiply-located universals which, if they exist, are compatible with Humean Supervenience.

So there are versions of endurance theory which are compatible with Humean Supervenience. Conversely, there are versions of perdurance theory which are incompatible with Humean Supervenience: there might be nonHumean relations between temporally-separated stages. Nevertheless, the suspicion lingers that endurance is somehow inherently in tension with Humean Supervenience, whilst perdurance is not. What could be the source of that tension?

Perhaps endurance is incompatible with the stronger thesis of Humean Determination, according to which all facts not only supervene upon but are determined or grounded by local qualities and spatiotemporal relations. Endurance theory has it that objects wholly present at different times can stand in the relation of identity: perhaps facts about identity cannot be determined by facts about continuity, even when they supervene upon such facts.

This claim is certainly not obviously true: on the contrary, those who advocate a principle of the identity of indiscernibles are often motivated by the view that identity facts must be determined by qualitative facts. Moreover, where qualitative continuity and identity are in step with one another (as with qualitatively-restricted endurance theory), it looks as if the continuity facts ground the identity facts, rather than vice versa. After all, the continuity may be explained causally, and this in turn explains the identity facts. Plenitudinous endurance theory also sits comfortably with Humean Determination: a line is the trajectory of an enduring object because it is a sequence of matter-filled points.

In any case, it's unlikely that Lewis would endorse Humean Determination. In 1983, he explicitly denied that brief stages are more fundamental than persisting objects:

When I say that persons are maximal R-interrelated aggregates of person-stages, I do not claim to be reducing 'constructs' to 'more basic entities'. (Since I do not intend a reduction to the basic, I am free to say without circularity that person-stages are R-interrelated aggregates of shorter person-stages.) Similarly, I think it is an informative necessary truth that trains are maximal aggregates of cars interrelated by the ancestral of the relation of being coupled together (count the locomotive as a special kind of car). But I do not think of this as a reduction to the basic. Whatever 'more basic' is supposed to mean, I don't think it means 'smaller'. (postscript B to 'Survival and Identity')

How else might we articulate the tension between endurance theory and Humean Supervenience? Isn't endurance theory somehow inherently non-local? After all, if endurance theory is true, there are relational arcs of identity connecting objects which exist at different times. But then also if perdurance theory is true, there are relational arcs of identity connecting objects which exist at different times – perduring, four-dimensional objects. No difference there.

Let's try again. If endurance theory is true, there are relational arcs of identity connecting objects which are *wholly present at different times*. But then if perdurance theory is true, there are relational arcs of non-identity connecting objects which are

wholly present at different times – momentary stages of perduring objects. To the extent that cross-temporal facts about identity are temporally nonlocal, so too are cross-temporal facts about non-identity, i.e. distinctness.

We have not yet located the conflict between endurance theory and Humean Supervenience. There are versions of endurance theory which, if they were true of the actual world, would be compatible with the contingent truth of Humean Supervenience, whilst there are versions of perdurance theory which are incompatible with Humean Supervenience. If anything violates the spirit of Humean Supervenience more generally, it may be the very distinction between perdurance and endurance theories itself, relying as it does on claims about identity and distinctness between objects wholly present at different times. Similar considerations apply to other debates involving extension and multi-location: perhaps the true Humean should resist both endurance and perdurance, both universals and tropes, and both extended and nonextended spatial simples.

Section 2: In Defence of Stages

Humean Supervenience in isolation does not entail perdurance theory, even for the actual world. Nevertheless, it features in Lewis's arguments for the theory. In his 'Postscript B' to 'Survival and Identity', Lewis provides an argument for the existence of person-stages, an argument which can easily be generalised to nonpersonal stages. He begins by explaining what a person-stage is supposed to be, for the benefit of those who profess not to understand the notion:

A person-stage is a physical object, just as a person is...it talks and walks and thinks, it has beliefs and desires, it has a size and shape and location...That is what I mean by a person-stage.

Lewis then argues that such stages exist:

First: it is possible that a person-stage might exist. Suppose it to appear out of thin air, then vanish again...

Second: it is possible that two person-stages might exist in succession, one right after the other but without overlap. Further, the qualities and location of the second at its appearance might exactly match those of the first at its disappearance. Here I rely on a *patchwork principle* for possibility: if it is

possible that X happen intrinsically in a spatiotemporal region, and if it is likewise possible that Y happen in a region, then also it is possible that both X and Y happen in two distinct but adjacent regions. There are no necessary incompatibilities between distinct existences. Anything can follow anything.

Third: extending the previous point, it is possible that there might be a world of stages that is exactly like our own world in its point-by-point distribution of intrinsic local qualities over space and time.

Fourth: further, such a world of stages might also be exactly like our own in its causal relations between local matters of particular fact.

Fifth: then such a world of stages would be exactly like our own simpliciter. There are no features of our world except those that supervene on the distribution of local qualities and their causal relations.

Sixth: then our own world is a world of stages...

Seventh: but persons exist too, and persons (in most cases) are not person-stages. They last too long. Yet persons, and person-stages, like tables and table-legs, do not occupy spatiotemporal regions twice over. That can only be because they are not distinct. They are part-identical: in other words, the person-stages are parts of the persons. ('PostScript to Survival and Identity', 76-77)

Humean Supervenience crops up twice in this argument. First, the 'patchwork principle' of step two, whilst not strictly a consequence of Humean Supervenience, is closely related to it. Then at step five, we infer from the existence of a perdurance world which is Humeanly-indiscernible from our own world to the conclusion that our own world is a perdurance world.

Endurance theorists might concede that Lewis has demonstrated the existence of sequences of person-stages, but argue that he has not demonstrated that the members of any given sequence are distinct from one another. Suppose that a spatiotemporal region is occupied by an object. Intrinsic facts about that region do not determine whether the object *also* occupies other spatiotemporal regions, either by being wholly present at more than one region, or simply by extending beyond the boundaries of the original region.

This is relevant at step two of the argument. If it is possible that a spatiotemporal region be occupied by a child, and if it is likewise possible that a spatiotemporal region be occupied by an adult, then the patchwork principle establishes that it is possible that two temporally-adjacent such regions be occupied by a child and an adult respectively. But it doesn't establish that this can occur without the child's being identical to the adult, in the sense endurance theorists intend. To secure this stronger conclusion, we need the further assumption that facts about whether two adjacent regions contain the same wholly present object are not fixed by qualitative facts about those regions. And this is the illegitimate assumption that all versions of endurance theory conflict with Humean Supervenience.

Even if the argument as a whole does not succeed, Lewis's final point about parts and wholes does tell against plenitudinous endurance theory, the view that there are as many enduring objects as there are time-like sequences of matter-filled points. This amounts to many, many enduring objects, one for each stage the perdurance theorist recognises, and one for each sum of such stages. It is a short simplifying step to accept that many of these objects are partially identical to one another – related as parts and wholes – rather than entirely distinct. But to accept this is to accept perdurance theory.

Plenitudinous endurance theory is, of course, ontologically vast. But so is Lewis's version of perdurance theory: he combines a liberal approach to positing stages with a mereological universalism which generates as many persisting objects as does plenitudinous endurance theory. For Lewis, however, the persisting objects are not to be accounted as extra ontology in any robust sense: it is double-counting to count the wholes as separate from all the parts. (This accounting method will be invoked again towards the end of this chapter.)

Section 3: Temporary Intrinsic

The principal and decisive objection against endurance, as an account of the persistence of ordinary things such as people or puddles, is the problem of temporary intrinsic. Persisting things change their intrinsic properties. For instance, shape: when I sit, I have a bent shape; when I stand, I have a straightened shape. Both shapes are temporary intrinsic properties; I have

them only some of the time. How is such change possible? (*On the Plurality of Worlds*, 203-4.)

Evidently, it is possible for Lewis to have different shapes by having each shape at a different time. The problem lies in understanding what it is about the passage of time which makes it possible for Lewis to be bent at one time and straight at another, given that he could not have both shapes simultaneously.

Perdurance theorists have an immediate solution to the problem: Lewis has a temporal part which is bent, and a distinct temporal part which is straight. It is the various temporal parts which exemplify these three-dimensional shapes, whilst Lewis himself – the perduring person – satisfies the corresponding temporally-relative predicates in virtue of the properties of his temporal parts. Lewis satisfies ‘is bent at t_1 ’ because he has a temporal part existing at t_1 which instantiates the property *bent*.

Endurance theorists are initially offered two alternative ‘solutions’ to the problem.

First solution: contrary to what we might think, shapes are not genuine intrinsic properties. They are disguised relations, which an enduring thing may bear to times. One and the same enduring thing may bear the bent-shape relation to some times, and the straight-shape relation to others. In itself, considered apart from its relations to other things, it has no shape at all. This is simply incredible, if we are speaking of the persistence of ordinary things. (It might do for the endurance of [intrinsically unchanging] entelechies or universals.) If we know what shape is, we know that it is a property, not a relation. (*On the Plurality of Worlds*, 204)

Lewis is tantalisingly brief here, and doesn’t add much when he discusses temporary intrinsics in later publications, though he asks ‘where have the monadic properties *bent* and *straight* gone? What have they to do with our new-found *bent-at* and *straight-at* relations...?’ (‘Tensing the Copula’ 4)

Second solution: the only intrinsic properties of a thing are those it has at the present moment. Other times are like false stories...This is a solution that rejects endurance; because it rejects persistence altogether. And it is even less credible than the first solution. In saying that there are no other times, as opposed to false representations thereof, it goes against what we all believe.

No man, unless it be at the moment of his execution, believes that he has no future; still less does anyone believe that he has no past. (*On the Plurality of Worlds*, 204)

Thus Lewis rejects presentism, the view that only what is present exists. Again, he is extremely brief, both here and elsewhere.

A third type of endurance-friendly ‘solution’ is missing from *On the Plurality of Worlds*, appears in a footnote crediting Sally Haslanger and Mark Johnston in ‘Rearrangement of Particles’, and finally gets Lewis’s full attention in ‘Tensing the Copula’. The idea is that instantiation – the having of properties – is relative to times: *having* is a three-place relation between objects, properties and times, so *bent* and *straight* retain their status as intrinsic properties. Lewis’s key objection here is that if we reify *having*, we invite Bradley’s regress: in virtue of what do Lewis, *bent* and t_1 mutually instantiate *having*? There are various versions of the third strategy, and Lewis’s various responses tell us plenty about his views on properties, instantiation, set-membership, and truthmaking. But endurance theorists would not need a third strategy were it not for Lewis’s implacable insistence that three-dimensional shape is an intrinsic property of ordinary objects. This insistence is the core of his argument against endurance theory.

What should we make of the argument from temporary intrinsics? To be frank, it seems uncharacteristically weak. As Lewis himself shows elsewhere, the intrinsic/extrinsic distinction is not straightforward, and intrinsic/relational even less so. Moreover, relations between concrete objects and times (whatever times are), are likely to differ from paradigm relations such as *taller than*: Moorean common sense does not specify whether three-dimensional shape is such a nonparadigmatic relation. And shape is an peculiar exemplar: an object’s shape may well be reducible to occupancy relations between the object and a region, together with the shape of the region itself. Finally, we shouldn’t prejudge the capacity of science to surprise us about whether three-dimensional shapes are relations between their possessors, on the one hand, and other concrete objects, frames of reference or fields on the other; the same goes for mass, colour, biological species, and so on. Yet Lewis calls the argument from temporary intrinsics the ‘principal and decisive objection against endurance’. We confront an exegetical puzzle.

Though Lewis himself does not make the connection explicit, I think that his remarks about temporary intrinsics can fruitfully be understood in the light of his Humean Supervenience. If ordinary temporary properties of objects were relations between objects and times, then those relations would not supervene upon the intrinsic properties of their relata: neither objects nor times would have enough genuinely intrinsic properties to generate the rich qualitative array of the world we see around us. Would these nonsupervenient relations be amongst those exceptionally permitted by Humean Supervenience? According to Lewis, these are ‘exactly the spatiotemporal relations: distance relations both spacelike and timelike, and perhaps also occupancy relations between point-sized things and spacetime points’ (‘Humean Supervenience Debugged’, 225 in *Papers in Metaphysics and Epistemology*).

The nonsupervenient relations in question would not be mere occupancy relations – very many objects all occupy, or exist at, any given time, yet on this view these objects all stand in different ‘qualitative’ relations to that particular time (conversely, a single object occupies many different times, yet stands in different ‘qualitative’ relations to those different times). Nor are the nonsupervenient relations comparable to distance relations, which relate points to points, or objects to objects, not objects to times; nor are they like relations of orientation, betweenness or other spatiotemporal relations. So ‘temporary intrinsics as relations’ would violate Humean Supervenience.

Moreover, this explains why Lewis doesn’t consider empirical objections to the view that three-dimensional shape is intrinsic: he does not aspire to defend Humean Supervenience against scientific challenges. ‘Really, what I uphold is not so much the truth of Humean supervenience as the *tenability* of it. If physics itself were to teach me that it is false, I wouldn’t grieve.’ (*Philosophical Papers volume II*, xi; cf ‘Humean Supervenience Debugged’ 226). Lewis acknowledges quantum entanglement as a scientific challenge to Humean Supervenience, and likewise he can accept empirical claims that, for example, shape or mass might be a nonsupervenient relation between objects and frames of reference, or between objects and fields. His commitment is to resisting ‘philosophical’ suggestions of nonsupervenient

relationality: attempts to impose an endurance-friendly account of temporary properties are philosophical in this sense, so Lewis must resist them.

Recognising ‘temporary intrinsics as relations’ as a violation of Humean Supervenience also explains Lewis’s insouciance about a relational feature of his own view. Given perdurance theory, the temporary properties of persisting objects are relations between those objects and their temporal parts. Lewis is bent at t_1 , but Lewis himself is four-dimensional, and does not instantiate three-dimensional shape properties like *bent*. He is bent at t_1 only in virtue of the relation between Lewis and his temporal part Lewis-at- t_1 , which does instantiate the intrinsic property *bent*. ‘Simply incredible’? Apparently not. The acceptability of this perdurantist relational view of change may be explained by its compatibility with Humean Supervenience: an object’s relations to its own parts are amongst its intrinsic features, and *bent* is an intrinsic property of the temporal part in question.

So Lewis’s objection to ‘temporary intrinsics as relations’ may be grounded in his Humean Supervenience. And temporary intrinsics may be the ultimate source of the conflict between endurance theory and Humean Supervenience, a conflict which we earlier struggled to locate. As we saw, neither qualitatively restricted nor plenitudinous endurance theory is committed to nonsupervenient relations between different points. Instead, nonsupervenience is a feature of the relations between an enduring object and the succession of points it occupies, given that the object changes in its properties over time, and given the temporary-intrinsics-as-relations account of such change. If this is a problem, it is a problem for all versions of endurance theory which permit widespread change, including plenitudinous and qualitatively-restricted endurance theory; the only possible solutions, in Lewis’s eyes, are to reject Humean Supervenience, or else to explore other options such as presentism, or the relativisation of instantiation.

There is a lot to be gained from viewing Lewis’s discussion of temporary intrinsics against the backdrop of his Humean Supervenience. But some unresolved issues remain. First, there is the simple point that Lewis himself does not express his concern about temporary intrinsics as arising from his Humean Supervenience. Second, there is the vehemence of Lewis’s rejection of temporary intrinsics as

relations – ‘simply incredible...if we know what shape is...’ – given his relatively modest goal of upholding the mere tenability of Humean Supervenience. This goal requires only a demonstration of the tenability of a Humean-friendly account of change (perdurantism theory, for example), not a decisive crushing of non-Humean accounts of change such as endurance theory. So I will stop short of claiming that this is what Lewis had in mind. Nevertheless, making this connection enables us to strengthen the uncharacteristically weak argument from temporary intrinsics, and to pin down the elusive tension between endurance theory and Humean Supervenience.

Section 4: Stages, or Sums of Stages?

Lewis’s treatment of temporary intrinsics in *On the Plurality of Worlds* forms part of a discussion of both persistence and its modal analogue, for which we have no neutral term. For Lewis, possible worlds are concrete realities; amongst the parts of nearby non-actual worlds are concrete objects which closely resemble actual objects.

Consider possible-Katherine, a non-actual thing which closely – but not exactly – resembles actual-Katherine. Lewis develops an argument from accidental intrinsics which parallels his argument from temporary intrinsics, concluding that no concrete object is wholly present in more than one world: actual-Katherine and possible-Katherine are not literally one and the same thing, because they do not share all their intrinsic properties. Thus the modal analogue of endurance theory is rejected.

So what am I, Katherine, this person? There are two options. Either I am actual-Katherine, or I am the sum of actual-Katherine and all the possible-Katherines, a transworld individual with a part in each of many worlds. Lewis accepts that both these objects exist: his mereological universalism (belief in arbitrary sums) guarantees the existence of the transworld object, given the existence of the various worldbound objects. But Lewis identifies ordinary objects with worldbound individuals:

Katherine is actual-Katherine. Actual objects have nonactual counterparts, and the ordinary features of those counterparts ground the truth or falsity of claims about the modal features of ordinary objects. I could have been a contender, because I have a counterpart, a possible Katherine, who is a contender in the world she is a part of.

Lewis often exploits analogies between time, space and modality (e.g. *Philosophical Papers Volume I*, xi), and he rejects both endurance theory and its modal analogue.

But the analogy is not complete: for Lewis, ordinary objects are worldbound but they are not timebound. Perdurantism is challenged by the stage theory of persistence, which identifies ordinary objects with brief stages instead of transtemporal sums of those stages; claims about the past and future of a given timebound stage are made true by counterpart stages existing at different times, just as modal claims about worldbound objects are made true by counterparts existing at different worlds. Stage theory is the temporal analogue of Lewis's modal counterpart theory. So why isn't Lewis a stage theorist?

Lewis doesn't tell us directly why he rejects stage theory. Instead, he assumes that ordinary objects are transtemporal sums of stages (once endurance theory has been defeated), then provides several reasons why, even given this assumption, we should accept that ordinary objects are worldbound. Identifying ordinary objects with transtemporal and transworld objects seems to be the default option, to be accepted in the absence of countervailing reasons, and Lewis provides such countervailing reasons only in the modal case. Nevertheless, examining these reasons provides some insight into the relative merits of perdurantism and stage theory.

Lewis's "simplest" point is that most of us do not believe in transworld individuals, since we do not believe in non-actual concrete worlds. So identifying ordinary objects with actual-world-bound individuals provides at least one small respect in which Lewis can conform to "common opinion" (*On the Plurality of Worlds* 220). He doesn't explore the temporal parallel, but there are both philosophers and normal people who deny the existence of the future: if this fact carries any argumentative weight at all, it favours stage theory.

A different argument seems to apply to people, but not to other ordinary objects. Lewis points out that person-stages have a special attitude towards 'their' other-time stages, but not towards their other-worldly counterparts. "To some extent, stages want to fulfil the remembered desires of earlier stages... To a greater extent, stages want to fulfil the foreseen desires of later stages: that is prudence." In contrast "[m]y this-worldly self has *no* tendency to make the purposes of its other-worldly counterparts its own" (*On the Plurality of Worlds* 219). This is because there is no

transworld causation, and because there is no contingency in the collective fortunes of one's counterparts.

It is clear that transtemporal relations amongst personal stages differ from transworld relations in the way Lewis indicates. But the discussion brings out a difficulty for perdurance theory, not an advantage. Stages instantiate ordinary temporary properties, psychological and otherwise: this is key to Lewis's account of temporary intrinsics, and it is also important to his treatment of self-locating beliefs in 'Attitudes De Dicto and De Se'. When you lie awake at night wondering what time it is, your stage wonders where *it* is located in time, which of many wakeful stages it is.

So, as Lewis admits:

...my view makes it fair to think of the desires as belonging in the first instance to my present stage, and derivatively to the persisting sum of many stages. And I agree that what my present stage wants is not to finish the book itself – it's a sensible stage so it knows that can't be expected. But that's not to say that it only cares what happens, never mind how. There is a middle ground. My present stage wants the book to be finished in the fulfilment of its present intentions – there's the egocentric part – and that will happen only if the proper sort of causal continuity binds together my present stage with the one that finishes the book. (*On the Plurality of Worlds* 126, footnote 8)

Unlike stage theorists, perdurance theorists must distinguish the instantiators of ordinary temporal properties – *desiring, wondering, being bent, being straight, being alive* – from the instantiators of sortal properties like *being a person*. This counts against perdurance theory.

Still, Lewis is right that stages' concern for 'their' other-timely stages is mediated by distinctive causal connections which are absent in the transworld case. The point generalises: we care more about actual goings-on, regardless of whether they affect us personally, than we do about other-worldly matters, even though all are supposed to be equally real, equally concrete. (*On the Plurality of Worlds* section 2.6) Yet the thin, acausal nature of transworld relations does not directly show that people or other ordinary objects are not transworld entities, just as the richness of transtemporal

relations does not directly show that ordinary objects are transtemporal; more work is needed to bring out the connection.

The detail of Lewis's argument that ordinary objects are worldbound is set out in *On the Plurality of Worlds* 218-20, but we get to the heart of it as early as 1971:

We may draw an analogy between the [transtemporal] relations among stages of persons...and the personal...[modal] counterpart relations among [persisting] persons...But the analogy is imperfect. The [transtemporal] relations are equivalence relations, at least for the most part and as a matter of contingent fact...[but modal] counterpart relations are not equivalence relations. ('Counterparts of Persons and their Bodies' 52)

Both temporal and modal counterpart relations are reflexive, and, let us assume, symmetric. But transtemporal counterpart relations are typically transitive, whilst modal counterpart relations are not. That is, there are cases in which a is a counterpart of b , b is a counterpart of c , and yet a is not a counterpart of c .

Like all similarity relations on a sufficiently variegated domain, [modal counterpart relations] fail to be transitive because chains of little differences add up to big differences. ('Counterparts of Persons and their Bodies' 52)

Lewis's argument has two main steps. First, he must show that modal counterpart relations are just similarity relations, unlike temporal counterpart relations. This establishes that modal counterpart relations are not transitive. Second, he must show that the nontransitivity of modal counterpart relations establishes that ordinary objects are worldbound, whilst the few exceptions to transitivity permitted by the temporal counterpart relations do *not* establish that ordinary objects are timebound stages. I will discuss these two main steps in turn.

In *On the Plurality of Worlds* (218), Lewis shows us that temporal counterpart relations are constrained in ways which do not apply to modal counterpart relations. First: the temporal stages of ordinary objects are bound together by causal relations. Mereological universalism commits Lewis to very many transtemporal sums, most of which are not causally coherent. But those sums of stages which correspond to ordinary objects – people and puddles – do exhibit causal continuity. In contrast,

modal counterpart relations cannot involve causal continuity, for there is no transworld causation.

Second: times are ordered by the *earlier-than* relation (setting aside the complications of special relativity), and this in turn orders the stages. Thus, when we consider whether some stages are stages of the same ordinary object, we can put them in temporal order, then ask whether their collective qualitative variation over time is incremental. In contrast, the possible worlds are not set out in a neat row, and so a collection of world-bound individuals cannot be given a unique ordering. This means that modal counterpart relations cannot involve facts about ‘the objects in-between’ two putative counterparts.

These two points together do not quite demonstrate that modal counterpart relations are just similarity relations, but they make this very plausible by ruling out two natural alternatives: after all, what else could modal counterpart relations be? So whether or not one object is a counterpart of another depends entirely on their degree of similarity in various respects, respects determined by the relevant counterpart relation. Since logical space is as variegated as can be, such similarity relations are not transitive.

This brings us to the second main step of Lewis’s argument. If we accept that modal counterpart relations are similarity relations, and thus intransitive, how does this show that ordinary objects are worldbound? Intransitivity clearly shows that counterpart relations are not identity relations, because identity is transitive. But we have already accepted that possible worlds do not share parts, that counterparts are not literally identical with one another. Moreover, intransitivity does not rule out the existence of transworld individuals; these are mandated by Lewis’s mereological universalism. Instead, intransitivity tells against the identification of transworld individuals with ordinary objects, by undermining the ordinary ways in which we count ordinary objects.

Suppose for the sake of argument that Katherine is a transworld individual. Exactly which transworld individual? Katherine is a maximal sum of worldbound individuals, including actual-Katherine, mutually linked by the ‘same person as’ counterpart

relation. But intransitivity ensures the existence of an enormous number of these maximal sums.

How so? Consider two personal counterparts of actual-Katherine, possible-Katherine- α and possible-Katherine- ω . Each differs from actual-Katherine in significant but different ways, and these significant differences add up, so the two are not personal counterparts of one another (they exemplify the intransitivity of counterparthood). Then actual-Katherine and possible-Katherine- α are both parts of one transworld individual, whilst actual-Katherine and possible-Katherine- ω are both parts of a second transworld individual. Each of these is an equally good candidate to be the transworld Katherine, with nothing to choose between them. And there are many more such candidate Katherines.

If people are transworld individuals, then we have a population crisis in the actual world. I seem to be alone in my office as I write this, but in fact a multitude of transworld people are here right now, typing these words. Moreover, proper names for people – and other ordinary objects – are massively ambiguous. To recover the ordinary verdict that there is just one person in my office right now, we must insist that we can count people by counting their worldbound components, or that we needn't count all the people, or that we count all the people but do not count them by identity. This choice of evils constitutes the case against identifying ordinary objects with transworld individuals.

So why the difference for transtemporal individuals? Temporal counterpart relations are richer than mere similarity, but the additional requirements of causal continuity and incremental change do not guarantee their transitivity. There are possible situations – familiar from science fiction – in which a present stage bears the temporal counterpart relation to each of two different stages existing at the same future time, stages which are not counterparts of one another: 'fission' is possible. Ditto for past times: 'fusion' is possible. If fission lies in my future, then there are two perduring people in my office right now, and we must explain away the commonsense verdict that there is just one person here.

But such violations of transitivity are rare or nonexistent in the actual world:

We get by because ordinary [temporal] cases are not pathological. But modality is different: pathology is everywhere...So the modal case will always, or almost always, give us the same choice of evils about how to count [ordinary objects] that the temporal case gives us only in connection with far-fetched stories. (*On the Plurality of Worlds* 219)

For contingent reasons, perdurance theory rarely if ever forces us to reconstrue the ordinary ways in which we count actual ordinary objects (people, at least). But stage theory goes one step better: even if ‘fission’ will occur, we do not need to choose between counting stages and counting ordinary objects, for stages just are ordinary objects. So concerns about intransitivity and counting favour stage theory over perdurance theory, just as in the modal case. But because temporal pathology is rare, this advantage is small, and may easily be outweighed if stage theory has other disadvantages. As I noted above, Lewis seems to take perdurance theory as the default view once endurance theory has been excluded, and this small advantage for stage theory may not be enough to undermine perdurance theory’s status as default.

The story doesn’t end there, however, for Lewis’s mereological universalism and attitude to indeterminacy means that he too must reconstrue the ways in which we count ordinary objects, even in nonpathological situations. Most ordinary objects have vague spatial and temporal boundaries: there are stages around the end of my life such that it is indeterminate whether those stages are temporal parts of me. For Lewis, this amounts to indeterminacy in which of many slightly-different precise objects satisfies ‘is a person’, which one bears my name.

Do you think you are one person? – No, there are many aggregates of temporal parts, differing just a little at the ends, with equal claim to count as persons, and equal claim to count as you. Are all those equally good claims good enough? If so, you are many. If not, you are none. Either way we get the wrong answer. For undeniably you are one. (‘Many But Almost One’ 165 in *Papers on Metaphysics and Epistemology*)

Ultimately, Lewis is untroubled by this multiplicity because the person-candidates overlap very significantly, sharing most of their parts: they are ‘partially identical’,

indeed 'almost identical'. (I mentioned a similar move in section 2: the ontological extravagance of perdurance theory is reckoned less serious than the ontological extravagance of plenitudinous endurance theory.) This reconciliation isn't available where multiplicity arises from fission or fusion cases, since the objects in such cases differ significantly from one another; all the more so in the universal pathology of the modal.

Not all of us share Lewis's preference for perdurance theory over stage theory, and of course not all of us follow him in rejecting endurance theory. But his treatment of this issue – which barely reaches the surface of his writing – is an especially beautiful example of Lewisian metaphysical and semantic views working to support one another. It illustrates the systematicity of his metaphysics, integrating the rejection of worldly indeterminacy and the quasi-identification of parts and wholes with headline issues about worlds, times and properties. And it offers a master-class in the art of recognising both the power and the limitations of analogy.