

SECOND-ORDER PREDICATION AND THE METAPHYSICS OF PROPERTIES¹

Andy Egan

Problems about the accidental properties of properties motivate us—force us, I think—not to identify properties with the sets of their instances. If we identify them instead with functions from worlds to extensions, we get a theory of properties that is neutral with respect to disputes over counterpart theory, and we avoid a problem for Lewis’s theory of events. Similar problems about the *temporary* properties of properties motivate us—though this time they probably don’t *force* us—to give up this theory as well, and to identify properties with functions from ⟨world, time⟩ pairs to extensions. Again, the replacement theory is neutral with respect to a metaphysical dispute that the old theory (arguably) forces us to take a stand on—the dispute over whether objects have temporal parts. It also allows us to give a smoother semantics for predication, to better accommodate our intuitions about which objects temporary properties are properties of, and to make temporally self-locating beliefs genuinely *self-locating*.

Introduction

If you’re doing systematic metaphysics, problems for, and revisions to, one part of your theory will tend to have ramifications elsewhere. This paper is about a problem for Lewis’s theory of properties. It’s also about what happens to the Lewisian system when we revise the theory of properties in order to avoid the problem. So I’ll be concerned to do two things in what follows: First, I’ll argue that Lewis’s theory of properties has a fatal problem with accounting for second-order properties—the properties of properties. Second, I’ll trace the ramifications of the revision through the rest of Lewis’s system.

I. The Problem

An attractive nominalist strategy is to identify properties with the sets of their instances—not their *actual* instances (the property of *being round* isn’t the property of *being blue*, even

¹ Thanks to Tyler Doggett, Liz Harman, Sarah McGrath, Jim John, Juan Comesaña, Doug Marshall, Tony Corsentino, Daniel Stoljar, Vann McGee, Alex Byrne, Bob Stalnaker, Steve Yablo, Sally Haslanger, Jeff King, Michael Glanzberg, Laurie Paul, Adam Elga, Michael Smith, John Hawthorne, Adam Sennett, Brian Weatherston, Karen Bennett, and most of all to Ned Hall for discussion, comments, criticisms, suggestions, objections, and guidance.

if all and only the round things are blue)—but all of their instances, wherever in logical space they might be.² This is the theory that Lewis endorses [1986a; 1999].³

It's a straightforward consequence of this theory that things that exist in more than one world can't have any of their properties accidentally. Here is why not:

Suppose (contra Lewis) that people aren't worldbound—that they exist in more than one world. Elmer is a philosopher, but he might have been a plumber instead. So he's only *accidentally* a philosopher. Though he's a philosopher in the actual world @, he's not a philosopher in some other possible world *w*. That is: in @, Elmer has the property, *being a philosopher*, while in *w*, Elmer lacks the property, *being a philosopher*. Suppose also (with Lewis, this time) that properties are sets of possibilia. Something has the property just in case it's a member of the set, and lacks the property just in case it's not a member of the set. So since Elmer is a philosopher in @, he must be a member of the set, *being a philosopher*. Since Elmer is not a philosopher in *w* (since, in *w*, he lacks the property, *being a philosopher*), Elmer must not be a member of the set, *being a philosopher*. So Elmer must both be and not be a member of the same set. Contradiction.

The reason why this kind of case *isn't* a problem for Lewis is also straightforward. Lewis takes philosophers to be worldbound, and to have their modal properties in virtue of the behaviour of their (distinct) *counterparts* in various other worlds [Lewis 1968; 1971; 1973: 39–43; 1986a: 192–263]. And while Elmer can't both be and not be a member of the same set, there's no problem about Elmer being a member of some set that not all of his counterparts are members of.

Now consider a parallel, though less familiar, case. Elmer has a favourite property. It's *being green*. But Elmer is fickle—he might have favoured some other property instead. So *being green* might not have been Elmer's favourite property. In fact, *being green* might not have been *anybody's* favourite property. So *being green* is somebody's favourite property, but it might not have been. Here's another way of saying the same thing: *being green* has the property, *being somebody's favourite property*, but only accidentally. This is a case of *second-order predication*—of attributing properties to properties. More specifically, it's a case of *contingent second-order predication*—of attributing *accidental* properties to properties. Since *being somebody's favourite property* is itself a property, it's the set of its (actual and possible) instances. Its instances are properties, so *being somebody's favourite property* is a set of properties.

² Sets, or classes? Well, Lewis *says* that he means sets [1986a: 50, n37]. But this runs into problems: *being a set*, for example, seems like a perfectly good property, but there isn't (on the standard ways of doing set theory) a set of all sets. So if properties have to be sets, *being a set* turns out not to be a property after all. (Classes run into other problems—*being non-self-membered* seems like a perfectly good property, too.) Since I won't be concerned with *these* problems, and nothing will hang on the set/class distinction in what follows, we needn't resolve this question here. (Thanks to Vann McGee for pressing me on this.)

³ There are a couple of extremely contentious bits of Lewisian metaphysics that I'm going to treat as if they were uncontroversial. The first of these is Lewis's modal realism. The other is that *some* sort of reductive nominalist theory of properties is right. The sorts of questions that are addressed in the remainder of the paper will still arise for many philosophers who reject one or both of these theses. First, while the questions only arise if you're *some* kind of modal realist, they don't depend on merely possible worlds being the big concreta that Lewis takes them to be. Second, while anti-nominalists won't want to *identify* properties with set-theoretical entities built out of possibilia, properties still need to be intimately associated with some such entities. Whatever properties are, they at least need to *determine*, for example, a class of all of their possible instances.

Let @ be the actual world, in which *being green* is Elmer's favourite property, and let *w* be a world in which Elmer (along with everyone else) has turned his attentions elsewhere. Since *being green* is somebody's favourite property in @, it must be a member of the set, *being somebody's favourite property*. Since *being green* is not anybody's favourite property in *w*, it must not be a member of the set, *being somebody's favourite property*. Contradiction. And since properties aren't worldbound and therefore don't have distinct counterparts in other worlds, counterpart theory can't help us here.⁴

So here's the argument: If properties are the sets of their instances, then properties can't have any of *their* properties accidentally. But properties plainly *can* have some of their properties accidentally—witness *being somebody's favourite property*, *playing the pain role*, and *being the semantic value of 'red'*, all of which are accidental properties of any property that has them. So properties aren't the sets of their instances.⁵

II. Objections and Responses

Isn't there some sort of paraphrase strategy available that'll make this problem go away? I don't think that there is. At the very least, there's not one that's worth the trouble, given that there's another solution available (presented in the next section) that gives a smooth treatment of the problem cases.

One thing that we might try is world-indexing all of the second-order properties.⁶ So there's no such property as *being somebody's favourite property*—instead there's a family of world-indexed properties, *being somebody's favourite property at w_1* , *being somebody's favourite property at w_2* , etc. These properties are had necessarily if at all, so we needn't say that properties have any of their properties accidentally. We can explain the *appearance* of having accidental properties by noting that *being green* has some, but not all, of these world indexed properties—a fact naturally expressed by saying that *being green* is somebody's favourite property in some worlds, but not in others.

This move fails because it assigns the wrong contents to sentences involving second-order predication. When I say, '*being green* is somebody's favourite property',⁷ two things

⁴ Mark Heller [1998] has proposed a theory on which properties *do* have distinct counterparts in the other worlds (in response to a very different problem). This sort of theory will, obviously, avoid the problem that I'm raising here.

⁵ Some more examples of accidental second-order properties: *being instantiated*, *being coinstantiated with being green*, and *being the subject of extended philosophical debate*. Some slightly fancier ones: the second-order properties you get by lambda abstraction from statements of contingent natural laws, and the second-order properties you get by dropping one of the quantifiers over properties in a Ramsey sentence, like the ones that (we hope) provide the characteristic functional roles of mental properties like *being in pain* (*being in pain* is the property G such that there are properties $F_1, F_2 \dots F_n$ such that ...).

⁶ Since properties are *arbitrary* sets of things, there won't actually be a clean distinction between first-order and second-order properties, since there will be properties that could be had by, for example, both *being green* and my desk. These properties actually needn't even be especially weird or unfamiliar; *being mentioned in this footnote* seems to do the trick. Still, there will be *some* properties that can only be had by particulars, and some that can only be had by properties—call these the *pure* first-order and second-order properties, respectively. Where it's important that the second-order properties not also be properties of particulars, read 'second-order property' as 'pure second-order property'.

⁷ Or even, '*being green* is Elmer's favourite property'. I focus on the case of 'quantified' properties like *being somebody's favourite property* not because the points only go through for these sorts of

have happened: I've attributed some property to *being green*, and I've said something (metaphysically) contingent. But suppose the property I've attributed to *being green* is a world-indexed property. Then I *haven't* said something contingent. What I've said must be either necessarily true or necessarily false. So since the contents of sentences like '*being green* is somebody's favourite property' are contingent, the second-order properties attributed in them can't be world-indexed.⁸

Another tempting strategy is based on the very plausible idea that the only accidental features of properties are which *relations* they stand in to other things. This suggests a way to avoid the problem—we could get the right sentences to come out true by exploiting the unproblematic contingent *relations* between properties and other things, while denying that they have any troublesome accidental *properties*.⁹

Tempting though it is, this strategy doesn't work, because the contingent relations *are* problematic—accepting the contingent relations also commits us to accepting the accidental properties.

It's non-negotiable that there are contingent facts about which relations properties stand in to other things (including other properties). If we're allowed to move from the contingent facts about which *relations* properties stand in to contingent facts about which *relational properties* they have, then the jig is up, since the Lewisian theory of properties can't allow that there are any contingent facts about the properties (relational or otherwise) of things that don't have distinct counterparts.

We plainly *can* move from the fact that Elmer is hunting Bugs (that Bugs stands in the *being hunted by* relation to Elmer) to the fact that Bugs has the relational property, *being hunted by Elmer* (and therefore that he has the property, *being hunted by somebody*). And if we can do that, we ought to be able to move from the fact that *being green* is Elmer's favourite property to the fact that *being green* has the relational property, *being Elmer's favourite property*, and therefore has the property, *being somebody's favourite property*. And since it's contingent that *being green* is anybody's favourite property, it's contingent that *being green* has the property, *being somebody's favourite property*. Now we have our counterexample to the claim that properties are the sets of their instances: a property that has a property, and has it contingently.¹⁰

⁷ *continued*

properties, but because showing that they go through for properties like *being Elmer's favourite property* requires a fair amount of house-to-house fighting about the details of counterpart theory.

⁸ Why would it be so bad if it turned out that, when I say '*being green* is somebody's favourite property', what I've said is necessarily true? Because if we say this, we get the facts about entailment wrong. The content of my utterance of, '*being green* is somebody's favourite property' ought to be entailed by the proposition that *being green* is Foghorn's favourite property and it ought to entail that somebody has a favourite property. It ought not to be entailed by the proposition that Foghorn is a rooster. If the world-indexer is right, the content of my actual utterance is necessary, and so is entailed by everything, and only entails other necessary truths. (Thanks to Sally Haslanger and Jeff King for discussion of this point.)

⁹ This move may look like a non-starter, because of the existence of contingent relations between properties (*being systematically coinstantiated with, being related in lawlike way L to, etc.*). Perhaps this really is a fatal problem with the proposal. It might be, though, that whenever two properties stand in a contingent relation to one another, it's because the properties stand in contingent relations to certain worldbound things. In any case, whether this is true or not, there is an independent, fatal problem for the proposal.

¹⁰ Here is almost the same argument, cast more formally: $[Rab]$ is true iff $[\lambda x(Rxb)a]$ is true. And $[\lambda x(Rxb)a]$ is true iff a has the relational property that's the semantic value of $[\lambda x(Rxb)]$. So if (a)

Notice that *any* story we tell in order to avoid attributing accidental properties to properties will require us to say that the move from relations to relational properties is illegitimate. We'll be forced to say, for example, that *being green* can stand in the *being the favourite property of* relation to Elmer without having the relational property, *being Elmer's favourite property*. I suspect that this is incoherent. But in any event, whether it's coherent or not, it's a desperate move.

Lewis says some things that suggest a different line of response. In *On the Plurality of Worlds*, Lewis says that '[a] universal can safely be part of many worlds because it hasn't any accidental intrinsics' [1986: 205, n6] and '[i]f indeed there are no accidental intrinsics to raise a problem, then overlap confined to the sharing of universals seems entirely innocent' [ibid.: 205].¹¹

It looks like the claim is that there's an important difference between accidental *intrinsic* properties and accidental *extrinsic* properties—there's a big problem if would-be trans-world objects (like universals) have accidental *intrinsic*s, no problem if their only accidental properties are extrinsic or relational.¹²

But the mere fact that the relevant properties aren't *intrinsic* doesn't seem like it can make a difference. The problem is that (a) there are things without distinct counterparts (namely properties) which plainly have (second-order) properties, and have them accidentally, and (b) Lewis's theory of properties cannot accommodate this fact. That the second-order properties are extrinsic is beside the point; the thing that's making the trouble is just that they're *properties*. When Lewis says that properties are the sets of their instances, he's offering a perfectly general theory of properties, not a theory that's only supposed to apply to some restricted subset of all the properties that there are. The claim that properties are the sets of their instances is supposed to apply to *all* properties, not just the intrinsic ones.

¹⁰ *continued*

there are some contingent sentences of the form [Rab] where one of the relata is a property, and (b) we can do the usual kind of lambda abstraction on sentences of the form [Rab], then properties can't be the sets of their instances. (Since we get some contingent facts about the properties of properties.) The properties we get by this procedure are ones like *being Elmer's favourite property*, rather than *being somebody's favourite property*, but (a) those properties will do just fine as counterexamples (see note 7), and (b) we can also get the 'quantified' properties if we want them by adding another innocent-seeming step to the argument, in which we existentially generalize before doing the lambda abstraction: [Rab] is true only if $[\exists y(Ray)]$ is true, which is true iff $[\lambda x[\exists y(Rxy)]a]$ is true. So in order for [Rab] to be true, a certain object must have the property expressed by $[\lambda x[\exists y(Rxy)]]$.

¹¹ Another relevant passage is [Lewis 1999: 11, n5]. In these passages, Lewis is talking about universals, which would be points of *overlap* between worlds—the very same thing would be part of more than one world, and not by having different parts in each. This is unlike Lewisian properties, which are *not* meant to be points of overlap. They do draw their members from many different worlds, and so it's natural to say that they have *parts* in various different worlds (especially for Lewis—see [1986b; 1990]). However, at least for first-order properties, a property's part in one world will be wholly distinct from its part in any other world. But this difference between properties and universals is not relevant to our concerns. What's important to generating the problem that Lewis is worried about for universals, and that I've claimed actually arises for Lewisian properties, isn't *overlap*, but the absence of distinct counterparts in different worlds. The problem is that universals, being present in more than one world, are always their own counterparts. So, it can't be that a universal is a member of some set and one of its otherworldly counterparts is not. This feature—the lack of distinct counterparts—is a feature that universals and Lewisian properties share.

¹² Stephen Yablo [1998] says some similar-sounding things.

In light of this, I think that Lewis's (very brief) discussion of accidental intrinsics is best understood not as a new line of argument, but as a gesture in the direction of one of the strategies discussed above.¹³

Summing up: we can't paraphrase away accidental second-order properties without saying implausible things about the connection between relations and relational properties and/or assigning the wrong contents to sentences involving second-order predication. We also can't appeal to the bare fact that the accidental properties of properties are all extrinsic in order to solve our problem. Fortunately there is another theory of properties right around the corner, in the same spirit as Lewis's theory, which handles second-order predication without incident.

III. The Replacement Theory

Perhaps surprisingly, there's no parallel problem for a theory that identifies properties with *functions* from possible worlds to extensions (where extensions are just sets of things).¹⁴

Here's how the functions account handles the fact that *being green* is somebody's favourite property, but only accidentally: At @, *being green* is somebody's favourite property. At *w*, it's not. The second-order property, *being somebody's favourite property*, is a function that, for each world taken as argument, delivers as value an extension, which will be a set of properties. Call the value of *being somebody's favourite property* taking @ as argument 'A', and call the value taking *w* as argument 'B'. Since *being green* is somebody's favourite property in @, *being green* must be a member of A. Since *being green* is not anybody's favourite property in *w*, *being green* must not be a member of B. No problem. All we get from this is the unsurprising result that $A \neq B$; that *being somebody's favourite property* delivers different extensions for @ and *w*.

So while there's a problem about contingent second-order predication if we take properties to be the sets of their possible instances, there's no problem if we take them to be functions from worlds to extensions. I conclude that Lewisians should take properties to be functions from worlds to extensions, rather than the sets of their possible instances.

This lets us avoid the problem about accounting for the accidental properties of properties while still telling a uniform story about first-order and second-order predication: something has a property F at a world *w* iff it, or its counterpart at *w*, is a member of the value of F taking *w* as argument. It has F essentially iff it has F at every world (or at

¹³ Lewis [1986a: 201, 205] actually says some things that pretty explicitly suggest at least the world-indexing strategy. Thanks to Jeff King for pointing this out to me, and for convincing me that there's a fair bit more to say about the appeal to intrinsicness (though unfortunately there's not space to say it here).

¹⁴ This is potentially surprising because it's tempting, if you're only thinking about the properties of worldbound things, to take the two theories to be notational variants. Of course, *Lewis* didn't take them to be mere notational variants. One of the problems about offering Lewis a properties-are-functions theory as a way to avoid the problem about contingent second-order predication is that he offers an argument *against* identifying properties with functions [Lewis 1986a]. We'll get to this in the next section.

every world at which it has a counterpart), and it has F accidentally iff it has F at @, and has the complementary property (it's in the antiextension of F) at some other world.¹⁵

We can still build properties out of antecedently familiar things, and we can still have plentiful, arbitrary, gerrymandered properties, and a distinction between the more and less natural properties. It's as easy to have arbitrary, gerrymandered functions as it is to have arbitrary, gerrymandered sets of things, so there's no difficulty in accommodating extremely unnatural properties. There's also no more of a problem singling out the natural properties if we take them to be functions than if we take them to be the sets of their instances.

So taking properties to be functions from worlds to extensions does away with the problem about contingent second-order predication, and preserves the benefits of Lewis's theory. So far, so good.

IV. Functions, Properties, and Relations

Despite these benefits, there's a worry about the acceptability of this substitute theory of properties. Lewis offers an argument that we ought *not* to take properties to be functions from worlds to extensions. Here I paraphrase his argument (the examples are the same as in the original):

It's tempting to say that there are properties that things don't have or lack *simpliciter*, but only relative to this or that other thing. Since I'm thirsty at some times and not at others, I don't just have (*simpliciter*) the property *being thirsty*. Instead I have the property *being thirsty* relative to some times, but not relative to others. Since the road is surfaced in some places but not in others, the road doesn't have (*simpliciter*) the property, *being surfaced*. Instead it has the property *being surfaced* relative to some locations, but not relative to others. Since Ted is the father of Fred, but not of Ed, he doesn't have (*simpliciter*) the property, *being a father*. Instead he has the property *being a father* relative to some people (like Fred) but not relative to others (like Ed). Similarly, since nine numbers the planets in the actual world but not in every possible world, nine doesn't have (*simpliciter*) the property, *numbering the planets*. Instead it has the property *numbering the planets* relative to some worlds, but not relative to others.

If you think that things don't have or lack their properties *simpliciter*, but only relative to this or that time, location, person, or world, then you should think that properties are functions from various kinds of things to extensions. You should think that *numbering the planets* is a function from worlds to extensions, *being thirsty* is a function from world-time pairs to extensions, *being a father* (on one reading) is a function from individuals to extensions, and *being surfaced* is a function from locations to extensions.

But we shouldn't think any of this, because we should think that things have or lack their properties *simpliciter*. A 'property' which is had only relative to this or that other thing is, whatever formal apparatus we use to describe it, not a *property* but a *relation*. So what the proponent of identifying properties with functions from things of one sort

¹⁵ The complication about having the complementary property in some world (rather than just failing to have the property in some world) is there in order to allow things whose existence is contingent to have some of their properties essentially.

or another to extensions is proposing is really *doing away* with genuine properties and leaving us with only relations in their place. But it's obvious that things *do* have properties—it isn't *all* relations. And so we should have a theory of properties that allows things to have or lack them *simpliciter*, rather than merely relative to this or that other thing.

[Lewis 1986a: 52–3]

At a first pass, this is an appealing argument. But its appeal is based on the assumption that functions from locations, people, world/time pairs, and worlds to extensions are all in the same boat, and will sink or swim together as candidates to be the properties. Since it's clear that functions from locations to extensions, or from people to extensions, *can't* be properties, it must be that they all sink. But the various candidates that Lewis mentions *aren't* all in the same boat. Not all functions are equally good candidates to be properties.

Lewis tells us, '[i]n order to say what a meaning *is*, we may first ask what a meaning *does*, and then find out what does that' [1983a: 193]. This isn't a special principle about meanings. We can use the same procedure for other kinds of things, too. What *properties* do (among other things, but first and foremost) is provide semantic values for predicates.¹⁶ And what the semantic values for predicates do is determine an extension at each world. (More carefully: they combine with the semantic values of names, quantifier phrases, etc. to yield propositions. If we take propositions to be functions from worlds to truth values, then the semantic values of predicates need to determine an extension for each world.)

If *that's* what properties do, then functions from worlds to extensions do it better than sets of instances. Such functions can supply the semantic values of predicates that apply to properties, while sets of instances can supply semantic values only for predicates that apply exclusively to worldbound individuals.

Functions from locations or people to extensions aren't good candidates to be the properties because they aren't good candidates to be the semantic values of predicates. They aren't good candidates to be the semantic values of predicates because they don't do the thing that the semantic value of a predicate needs to do: determine an extension at each world. Functions from *worlds* to extensions are, unsurprisingly, *ideally* suited to the role of determining an extension at each world. So they're very good candidates to be the properties. It's not *being a function* that makes functions from locations, people, etc. to extensions ineligible to be properties. It's being the wrong *kind* of function.¹⁷

V. Some Consequences

Counterpart Theory

If we accept the revised theory of properties, we lose one of the arguments for counterpart theory. One reason to be a counterpart theorist is to avoid a problem about contingent

¹⁶ This is Lewis's view [1999: 16–18], and it seems to me that it must be right.

¹⁷ What about functions from (world, time) pairs to extensions? How good are they as candidates to be the properties? We'll come back to this in sections VI and VII.

first-order predication. If we take properties to be the sets of their actual and possible instances, then we have to be counterpart theorists, since otherwise the modally extended tables, chairs, and philosophers couldn't have any accidental properties. Once we take properties to be functions from worlds to extensions, we no longer have this reason to adopt a counterpart theory instead of a theory of transworld individuals.

Of course, there are other reasons to be a counterpart theorist.¹⁸ So the revision of Lewisian metaphysics I've suggested certainly doesn't require us to abandon counterpart theory. The surviving motivations—the implausibility of ordinary objects being only partly actual (or being points of overlap between worlds), scepticism about whether there are any deep facts about the essences of things, independent of our concerns or the particular ways in which we refer to them, a desire to identify statues with the clay that constitutes them, or persons with their bodies—seem to be the stronger, and the more philosophically interesting, motivations anyway. But still, it's an interesting fact that one of counterpart theory's selling points doesn't survive the revision. If the revised Lewisian is going to be a counterpart theorist, it won't be because she's compelled to by her theory of properties. Instead it will be due to some slightly more elaborate, and less obviously conclusive, argument about essentialism, the implausibility of ordinary objects being that widely scattered, etc.

Events

One place where the revision makes things *better* for Lewis's metaphysics is in the theory of events. Lewis [1986b] identifies events with properties of spacetime regions. It should come as no surprise that this view has trouble accounting for the fact that *events* have some of their properties accidentally. The problem is the one you'd expect. Lewis identifies events with properties of spatiotemporal regions. So, given his theory of properties, he identifies them with sets of possible regions. But this sort of modally scattered set isn't the kind of thing that has different counterparts in different worlds, and so, if the only story about having properties accidentally is the counterpart account that applies to ordinary world-bound things, events can't have any accidental properties. But of course they *do* have accidental properties, so there's some fancy footwork to be done in order to make the theory plausible.

Much of 'Events' is spent in an attempt to give an account of what it is for an event to have a property accidentally—more carefully, of what it is that makes apparent attributions of accidental properties to events true. As Lewis leaves it at the end of the paper, the account is remarkably complex and counterintuitive, and, as Lewis notes, incomplete—there are still accidental properties of events that are left unaccounted for. This is one of the major blemishes on Lewis's theory of events. One important thing that a theory of events ought to do is account for such plain facts as that yesterday's football game involved Brett Favre (featured 700 yards of total offense, came down to the wire, etc.), but might not have. It's easy to think that a theory of events that needs to take such desperate measures in order to account for this sort of fact can't possibly be right.

¹⁸ Some of them can be found in Lewis [1971; 1986a].

We can now see that the difficulty about accidental properties that Lewis was struggling with in ‘Events’ is an instance of a much more general problem with his theory of properties. And since we have a solution to the general problem about properties, we also have a solution to the specific problem about events. None of the elaborate manoeuvres in ‘Events’ are necessary if properties are functions from worlds to extensions. So we can keep the core of Lewis’s theory of events—that events are properties of regions of space-time—without buying ourselves a big problem about how events can have some of their properties accidentally.¹⁹ This makes Lewis’s theory of events (or at least, a Lewisian theory of events) substantially more attractive.

Let’s take stock of what’s happened so far. I’ve argued that it won’t do to identify properties with the sets of their instances, because if we do, then we can’t account for the fact that properties have some of their properties accidentally. I’ve also argued that the way to avoid this problem is to identify properties with functions from worlds to extensions. Accepting this revision has two interesting consequences: first, we lose one of the arguments for counterpart theory; and second, we avoid the biggest problem for Lewis’s theory of events. The rest of the paper is about an argument that we need to make a further revision—that we actually ought to identify properties with functions from ⟨world, time⟩ pairs to extensions—and the consequences of accepting it.

VI. Worlds Enough, or Times?

Functions from people, locations, etc. to extensions are bad candidates to be the properties. Functions from worlds to extensions are better candidates. What about functions from ⟨world, time⟩ pairs to extensions? There’s an argument for identifying properties with functions that take ⟨world, time⟩ pairs rather than just worlds as arguments, parallel to the argument in section I.

Here is a way that taking properties to be functions from *worlds* to extensions could get us in trouble: At a certain time t_1 , Sylvester is sitting, and therefore has a bent shape. But Sylvester isn’t *always* sitting. At some other time t_2 , Sylvester is standing on a rickety stepladder trying to reach Tweety’s cage. So at t_1 , Sylvester is bent, and at t_2 Sylvester is not bent. That is: At t_1 , Sylvester has the property, *being bent*, and at t_2 he lacks the property, *being bent*. Suppose properties are functions from worlds to extensions. Then *being bent* is one such function—call it ‘F’. Then, since Sylvester is bent at t_1 , he must be a member of

¹⁹ Note for aficionados of events: There are actually two problems for Lewis’s theory about the properties of events, and the revised theory of properties that I’ve suggested only solves one of them. (Though I think it’s the more serious one.) Lewis wants properties of events, in the end, to be explained in terms of the (more or less intrinsic, more or less natural) properties of the possible regions where the events occur. There shouldn’t be any ‘free floating’ properties of events. So one problem is: what’s the relation between the properties of events and the properties of the (actual and possible) regions in which they occur? The revision to the theory of properties is no help at all here. The other problem, though, is: can the Lewisian allow that, *however* events come by their properties, they have some of them accidentally? If the answer to this is ‘no’, then the theory is dead in the water. If the answer is, ‘yes, but only by saying a bunch of really elaborate, counter-intuitive stuff’, then the theory is at least in serious trouble. And *this* problem goes away when we say that properties are functions from worlds to extensions, rather than the sets of their instances.

F(@). And since Sylvester is not bent at t_2 , he must *not* be a member of F(@). Contradiction.

If this argument sounds familiar, that's because it is. It's essentially Lewis's argument from temporary intrinsics for the existence of temporal parts. (Rather, it's essentially the first part of that argument—the rest is ruling out the other candidate solutions to the problem.)

Lewis, of course, avoids this problem by adopting a metaphysic of temporal parts [Lewis 1986a]. For Sylvester to be bent at a time t is for Sylvester's temporal part at t to be bent—that is, for Sylvester's temporal part at t to be a member of F(@)—and there is no problem about some of Sylvester's temporal parts being members of F(@) while others are not.²⁰ This would-be problem about Sylvester's *temporary* properties, and its solution, are exactly parallel with the would-be problem about Elmer's *accidental* properties, and its solution. Different times play the role of different worlds in the attempt to make trouble, and temporal parts play the role of otherworldly counterparts in dissolving the problem.

In the earlier argument, we were still able to make trouble by finding a sort of entity that *had* accidental properties, but *lacked* (distinct) counterparts. Here again, we can revive the problem by finding a sort of entity that has temporary properties, but lacks temporal parts. In fact, the same kind of entity will do the trick.

At t_1 , Sylvester is seated (and therefore bent) and hungry (he's hatching a devious plan to catch and eat Tweety). At t_2 , after his plan has come to fruition, Sylvester is stretching (and therefore straight) and full. Then the property, *being bent*, is coinstantiated with *being hungry* at t_1 (because Sylvester is bent and hungry), but not at t_2 . (Suppose that nobody else is both bent and hungry at t_2 , either.) That is: *being bent* has, at t_1 , the property, *being coinstantiated with being hungry*, but lacks it at t_2 . Suppose that the property, *being coinstantiated with being hungry*, is a function from worlds to extensions—call it 'G'. Since, at t_1 , *being bent* is coinstantiated with *being hungry*, *being bent* must be a member of G(@). But since, at t_2 , it's not coinstantiated with *being hungry* (since at t_2 , nobody is both bent and hungry) it must not be a member of G(@). Contradiction.

Appealing to temporal parts won't help us in this case, since properties don't have temporal parts.²¹ So we really do have a problem.

Just as the problem here mirrors our original problem, the solution here mirrors our original solution. The problem goes away if we take properties to be functions from (world, time) pairs to extensions.

This allows the extensions assigned to properties to vary with times, not just with worlds. The property, *being coinstantiated with being hungry*, can return an extension that

²⁰ A complication: Sylvester has lots of different temporal parts that exist at t . (The one that starts at t and ends three seconds later, the one that starts five minutes before t and ends five minutes after, the one that starts twelve seconds after Sylvester's birth and ends at his death, etc.) Some of these are also bent at some times and not at others. So for 'Sylvester's temporal part at t ', we should read, 'Sylvester's minimal temporal part at t '; the relevant temporal parts are probably the instantaneous slices. (There's a worry about whether or not these slices are really the right kinds of things to have the relevant properties, but we won't worry about this for now.)

²¹ Is this so clear? Well, it's hard to see what the temporal parts of properties would be. The natural candidates to be the *parts* of properties, construed as sets of one sort or another (such as sets of instances, or sets of (world, extension) pairs), are their subsets. But properties won't, in general, have subsets that are well suited to play the role of temporal parts. (Especially if we've been convinced by the first part of this paper, and we're taking properties to be functions from worlds to extensions rather than sets of instances.)

includes *being bent* taking $\langle @, t_1 \rangle$ as argument, but return an extension that doesn't include *being bent* when it takes $\langle @, t_2 \rangle$ as argument. Problem solved. Properties are functions from $\langle \text{world, time} \rangle$ pairs to extensions.

It would be nice if this argument were conclusive. Unfortunately it's not. One of the unsuccessful objections to the original argument has a parallel here that isn't subject to the same criticism. Time-relativizing the second-order properties, it turns out, fares better than world-relativizing.

One of the responses discussed in section II was world-indexing second-order properties. We deny that there's any such property as *being somebody's favourite property*, and replace it with a bunch of world-indexed properties such as *being somebody's favourite property in w_{453}* . These properties are had necessarily if at all, so we're left without any troublesome accidental second-order properties.

This move failed because it assigned the wrong contents to sentences like '*being green* is somebody's favourite property'. The parallel time-indexing move seems to fare better.

We can avoid the problem about the *temporary* properties of properties by denying that there are any. Time-index all of the temporary second-order properties: there's no such property as *being coinstantiated with being hungry*—instead there is a family of time-indexed properties, such as *being coinstantiated with being hungry at t_1* . In this case, though, there's no obvious problem about assigning the wrong contents to sentences. The contents assigned to (particular utterances of) sentences like '*being bent* is coinstantiated with *being hungry*' on this view are perfectly respectable—indeed, intuitively exactly right.²² So we can resist the argument for taking properties to be functions from $\langle \text{world, time} \rangle$ pairs to extensions by time-relativizing all of the (apparent) temporary second-order properties.²³

There are some worries about doing away with temporary properties in favour of their time-relativized cousins. One is that the resulting picture of the world looks suspiciously like one in which there's not any change—there aren't any properties that things have at some times, but lack at others. (This is also a worry about the temporal parts view—really, for any view that doesn't, somehow or other, relativize property instantiation to times.)²⁴

There's also another semantically driven objection to time-relativizing properties. It would be nice if the semantic values of (at least most) predicates were constant across different contexts. This won't be so if the semantic values of predicates are time-indexed properties. Predicates like 'is red' will need to express different time-indexed properties on different occasions of use.

These objections aren't conclusive. I think that the cost/benefit assessment will, in the end, favour $\langle \text{world, time} \rangle$ functions, but the argument here certainly isn't as

²² A first-pass semantics: if Daffy says, '*being bent* is coinstantiated with *being hungry*' at t , then that occurrence of 'is coinstantiated with *being hungry*' expresses the time-indexed second-order property, *being coinstantiated with being hungry at t* , and so the sentence expresses the proposition that's true in all and only the possible worlds in which *being bent* is coinstantiated with *being hungry* at t —that is, all and only the worlds in which, at t , something is both bent and hungry. This sounds like an extremely good candidate to be the content of Daffy's utterance.

²³ Of course there are temporal analogues of the modal notions of contingency and necessity, but it's much less controversial (perhaps not controversial at all) to say that the propositions expressed by the relevant sentences are all *timelessly* true than it is to say that they're all *necessarily* true. At least part of the reason for this is that in the temporal case, there's no parallel to the problem about entailment discussed in note 8.

²⁴ This sort of worry appears in Haslinger [1989a].

open-and-shut as the case for the first revision. Suppose that you agree with me, and we accept the further revision. What happens to the rest of Lewis's system?

VII. Consequences of the Second Revision

The Puzzle of Change and the Argument from Temporary Intrinsic

Lewis argues that the best way to account for the phenomenon of (intrinsic) change—the fact that, for example, Sylvester can be bent at t_1 and not bent at t_2 —is to adopt a metaphysic of temporal parts. Sylvester is bent at t_1 because his temporal part at t_1 is bent, and he's not bent at t_2 because his temporal part at t_2 is not bent. And while Sylvester can't both have and not have the same property, there's no problem about some of his parts having properties that other parts don't. This, together with arguments against competing solutions, is the argument from temporary intrinsics.

Three things fall out of the discussion in the last section:

First, adopting a metaphysic of temporal parts is not a *general* solution to the problem of change. It's a solution that works for *particulars*, but not for things like properties, which don't have temporal parts.

Second, if we avoid the problem of change as it arises for second-order properties by taking properties to be functions from $\langle \text{world, time} \rangle$ pairs to extensions, there's no need to adopt a metaphysic of temporal parts for particulars. Properties such as *being bent* (just like properties such as *being coinstantiated with being hungry*) can return different extensions for the same world at different times. So we don't need to appeal to temporal parts to explain how Sylvester can be bent at some times and not at others, and the argument from temporary intrinsics loses its force.

Finally, while across-the-board time-indexing may be able to solve the second-order problem of change, it *also* defangs the argument from temporary intrinsics. If we time-index shape properties like *being bent*, then there's no problem about Sylvester (the cat himself, not just his parts) having some, but not all, of the various *being bent at t* properties. And so there's no need to invoke temporal parts in order to solve the problem of change.

In order for the argument from temporary intrinsics to be persuasive, the Lewisian needs a non-temporal-parts solution to the problem of temporary second-order properties that doesn't generalize. If there's no story about temporary second-order properties, then we haven't solved the problem of change. If the story about second-order properties generalizes, then we'll already have solved the problem of temporary first-order properties without any appeal to temporal parts.

I think that the role of *intrinsicness* in the argument from temporary intrinsics is to rule out as general solutions certain strategies (like time-relativizing) that seem to work for extrinsic properties. So it must be that, while we can, for example, time-index all of the extrinsic properties, time-indexing intrinsic properties like *being bent* is unacceptable.

One reason why we might think this (not Lewis's reason—see below) is that time-indexed properties are all *relational*—they're properties that we get from a relation (between objects and times) by fixing one of the relata (the time). If relational properties couldn't be intrinsic,

then we couldn't say that shape properties were time-indexed, since shape properties are clearly intrinsic. This would give us a way to accept the time-indexing solution for temporary second-order properties (which are all extrinsic), but to deny that it generalizes—we can't apply it across the board, because we'd be left without any intrinsic properties.

This argument is not persuasive. The claim that no relational property can be intrinsic has whatever force it has only because it's so clear that relational properties where the other relatum is something like a person, a location, or a fire hydrant can't be intrinsic. It's a hasty overgeneralization from these kinds of examples that leads us to think that *no* relational property could be intrinsic. (Lewis himself gives a counterexample: if Platonism is true, then *participating in the form of Squareness* is, though relational, still perfectly intrinsic [Lewis 2002: 3].)

Lewis's official reason for rejecting time-indexing (and for rejecting views according to which properties are relations to times) is that it does away with genuine monadic intrinsic properties like *being bent*. None of the time-indexed properties can be identified with *being bent*, and the *bent-at* relation that holds between objects and times is obviously not monadic. So it looks as if *being bent* has been left out of the picture. And this is very bad [ibid.: 4].

Notice that this objection hinges on *monadicity*, not on *intrinsicness*. It's not clear that there's any principled reason for thinking that it's more objectionable to leave the intrinsic *being bent* out of the picture than it is to leave out the extrinsic, but equally monadic (and equally necessary for systematic semantics) *being hunted by Elmer*, or *being somebody's favourite property*. If all monadic properties are on the same footing, then time-indexing won't work as the conservative Lewisian's non-generalizable solution to the second-order problem. Either Lewis's objection works, in which case time-indexing isn't a solution to the second-order problem, or it doesn't, in which case the time-indexing move *does* generalize.

Noticing that there's a problem about temporary second-order properties puts the argument from temporary intrinsics in a bad spot. For the argument to have force, we must have a principled reason for giving a disunified theory of change, on which we give one account of the temporary properties of things with temporal parts and another for the temporary properties of things (like properties) that lack them. (Or maybe one account for intrinsic properties and another for extrinsic ones.) There doesn't seem to be any such reason. So we ought not to be moved by the argument from temporary intrinsics.

This isn't really a consequence of taking properties to be functions from ⟨world, time⟩ pairs to extensions. It's a consequence of the fact that we need to give an account of change in things that don't have temporal parts, and there's no good reason to think that the account we give there won't generalize, and obviate the need to posit temporal parts in order to solve the problem of change in things like people, tables, and cats that plausibly *do* have temporal parts.

Still, the argument from temporary intrinsics isn't the only argument for temporal parts—there are a number of other arguments waiting in the wings [Lewis 1983b; Cartwright 1975; Balashov 1999; and especially Sider 2001]. Once we've adopted the second revision to the theory of properties, we don't need to appeal to temporal parts in order to solve the puzzle of change. But it's not as if the revision forces us to become endurantists. While the new theory of change is clearly *compatible* with endurantism, it's compatible with perdurance theories as well. What happens when we accept the revised theory of properties is that the problem of temporary intrinsics ceases to be an argument one way or the other.

Proper Subjects of Predication

We have a strong intuition that Sylvester's shape properties are *intrinsic*. One of the benefits of a four-dimensionalist ontology of temporal parts was supposed to be that it allowed us to respect this intuition in a way that competing theories did not.²⁵ There is an equally strong intuition that Sylvester's shape properties are intrinsic properties of *Sylvester*. The temporal parts view does not allow us to respect *this* intuition. It's Sylvester-stages, not *Sylvester*, the temporally extended continuant cat, that have the property *being bent*. The best that the temporal parts solution can do for Sylvester is to give him the time-indexed property of *being bent at t* in virtue of the shape of his t-stage. Considered as a whole, Sylvester doesn't have any particular shape (except maybe something describable in 4-d geometrical terms).

This fact has been used to object to temporal parts theories [Haslanger 1989a; 1989b]. One benefit of the proposed revision to the theory of properties is that we can avoid this problem. It's Sylvester, and not merely his parts, that has shape properties like *being bent*.

Semantics

Solving the problem of change by attributing temporary properties to temporal parts makes for an ugly semantics. Subject-predicate sentences will attribute properties sometimes to continuants, sometimes to their stages. The semantics will be particularly ugly if we, for example, time-index all of the extrinsic properties (alternatively, all of the second-order properties) and not the intrinsic (alternatively, the pure first-order) properties.

One good consequence of the revision is that it allows us to give a substantially cleaner semantics for predication. Suppose we solve the problem of temporary (first-order) properties by saying that properties like *being bent* are really properties of person-stages and cat-stages, rather than properties of temporally extended persons and cats. We still need to give an account of the cases where the properties being attributed are clearly properties of persons, not properties of stages. Doing this will require us to give a more complicated and disunified semantics than we'd like.

Here's one way it could go: Names like 'Sylvester' (and probably noun phrases in general) are ambiguous between continuant objects and their stages. When I say, 'Sylvester is bent', 'Sylvester' refers not to the whole cat, Sylvester, but to a Sylvester-stage. When I say, 'Sylvester is a cat', 'Sylvester' refers to Sylvester, the temporally extended continuant cat. When I say, 'Sylvester is a bent cat', 'Sylvester is a cat who is bent', or 'Sylvester is both bent and a cat', something fancy is going on. (For example, maybe the predicates are also ambiguous.)

Whatever the story is, it's going to be messier than we'd like.²⁶ It would be nice if people's (and cats') names reliably referred to the relevant people and cats. It would be nice

²⁵ It's not clear that this is really so—it's not clear why competing theories like time-indexing, or taking properties to be (world, time) functions, have any problem about intrinsicness. But leave this aside.

²⁶ How problematic this is, and how unique this problem is to temporal parts, is an interesting question. Chomsky [2000] discusses some similar examples that are independent of what we think about four-dimensionalism. We say things like, 'the tattered book on the table has been on the

to keep the systematic ambiguity of predicates to a minimum. And it would be nice if the attribution of temporary properties worked the same way across the board—not by sometimes attributing time-indexed properties to whole objects, and sometimes attributing unindexed properties to stages. If we take properties to be functions from ⟨world, time⟩ pairs to extensions, we don't need to make either noun phrases or predicates ambiguous in order to deal with the problem of change, and we can have a uniform story—we never need to do any time-indexing, and we can always attribute the properties to the whole object.

One seemingly less attractive consequence is that properties cease to be the kind of thing that, given an object, determines a possible-worlds proposition. On this account, properties will combine with objects to determine functions from ⟨world, time⟩ pairs to truth values, not functions from worlds to truth values. So when we get the semantic values of subject-predicate sentences by composing the semantic values of the subject expression and the predicate expression, we won't get possible-worlds propositions. We'll get *tensed* propositions—functions from ⟨world, time⟩ pairs to extensions—things whose truth values can vary at different times within the same world.

It's not clear what we should think of this consequence. If we were convinced that the semantic values of sentences had to be possible-worlds propositions, then it would be unwelcome. But it shouldn't be *that* unwelcome, for two reasons. First, it's not so clear that the semantic values of sentences have to be possible-worlds propositions rather than tensed propositions. Second, even if the contents of *most* sentences turn out to vary across times, we'll still have some eternal sentences whose truth values aren't affected by the time.²⁷

*Attitudes De Se*²⁸

Lewis [1979] argues that attitudes like belief and desire are best thought of, not as attitudes toward *propositions*, but toward *properties*. The reason for this is that propositions—thought of as functions from worlds to truth values—won't do as the contents of *self-locating* beliefs.

Suppose Pepe is lost in the library. This might be because he doesn't know which world is actual—because, for example, he doesn't know what the actual floor plan of the library is. But Pepe could still be lost, even if he knew everything there is to know about which world is actual. He could know the complete floor plan of the library, and even the location of every person, cat, and skunk in the library, and still not know where *he* is, because he could still fail to know which, of all of the creatures in the library, is *him*. (It helps, to make this sort of ignorance compatible with full knowledge of which world is actual, to suppose that it's a very boring library full of amnesiacs, all having experiences indiscernible from Pepe's.)²⁹

²⁶ *continued*

best-sellers list for twelve weeks now'. This requires some sort of special treatment, since only book-tokens are tattered, and only book-types make the best-seller list. (Though Chomsky puts these examples to a different use.)

²⁷ It's also not so clear that this view of the semantic values of predicates really commits us to any particular view about the semantic values of sentences. If we still want sentences to be associated with possible-worlds propositions, we can adjust the rest of our semantic (and maybe syntactic) theory in order to get some other constituent of the sentence to do the extra work.

²⁸ Thanks to Robert Stalnaker for pointing out this consequence.

²⁹ Well, a library that's boring apart from the fact that it's full of amnesiac skunks.

The same concerns arise for beliefs that, for example, my pants are on fire. I can believe that *my* pants are on fire without believing that Egan's pants are on fire, and I can hope that someone turns a fire extinguisher on me right now without hoping that someone turns a fire extinguisher on Egan at 5:41pm. There also seems to be something that we both believe when I believe that *my* pants are on fire, and you believe that *your* pants are on fire. These facts are not easily accommodated on a view that takes attitudes like belief and desire to be attitudes toward possible-worlds propositions. Knowing which possible-worlds propositions are true can, at best, pick out a unique world as the one we inhabit. And like they say in the Bond movie, sometimes—when, for example, we want to know who or where we are—the world is not enough. We also need to know our place within it.

Lewis [1979] suggests that we treat self-locating attitudes as attitudes toward properties, rather than toward propositions. So (on Lewis's theory of properties) when I believe that my pants are on fire, I believe that I am one of the possible individuals with burning pants. My attitude is toward a property—*having burning pants*, and the content of the belief is that I am one of the individuals who instantiates the property.³⁰ We get essentially the same results if we accept the first revision and take properties to be functions from worlds to individuals.³¹

There's something very appealing about the treatment of self-locating beliefs as self-ascription of properties. But this strategy runs into trouble with *temporally* self-locating beliefs.

Lewis can't say that, when we have beliefs about what time it is, we're self-ascribing some property. All but the most unfortunate people are temporally extended. So our attitudes toward properties (considered as sets of individuals) will single out sets of temporally extended things—four-dimensional spacetime worms. But singling out a worm doesn't tell us what time it is. Though they don't say it in the Bond movie, sometimes—when, for example, we want to know what time it is—the worm is not enough.

If we say that properties are either the sets of their instances or functions from worlds to extensions, we can't treat my beliefs about what time it is as *self*-ascriptions of properties. No property of *Egan*—the continuant person—is going to make the relevant distinctions between different times. The most natural thing to say, probably, is that it's not *people* that have temporally self-locating beliefs, but their temporal parts. It's not *me*, but my present temporal part, that thinks that it's 5:47 (by thinking that it's one of the 5:47 person-stages).

³⁰ As opposed to my attitude being toward a proposition—that *Egan has burning pants*—and the content of the belief being that I am in one of the worlds in which the proposition is true.

³¹ It's a bit trickier to state just what's going on, though. The best way to see it is this: Functions f from worlds to extensions are in one-one correspondence with sets of $\langle w, i \rangle$ pairs (where w is a world and i an individual) such that i is a member of $f(w)$. So when we take self-locating beliefs, desires, etc. to be attitudes toward properties, we can think of them as attitudes toward, not sets of worlds, but sets of \langle world, individual \rangle pairs. If individuals are worldbound, then the worlds won't be doing any work—sets of individuals would do just as well. But this way of treating the contents of self-locating beliefs is compatible with (though it doesn't mandate) taking individuals to be present in more than one world. Notice that this mirrors proposals that treat the contents of self-locating beliefs as sets of *centred worlds*. Some philosophers also include a specification of a *time* in their centred worlds—we'll get to this in a moment.

But this sounds strange. It's certainly surprising to learn that people never have beliefs about the time.³²

It would be nice not to have to say this, and we don't have to if we accept the second revision. If properties are functions from $\langle \text{world}, \text{time} \rangle$ pairs to extensions, then people really can have temporally self-locating beliefs by self-ascribing properties.³³ So we get a nicer story about temporally self-locating beliefs if we accept the second revision: beliefs about what time it is can be treated as self-attributions of properties by people, not just by their stages.

Conclusion

Problems about accidental second-order predication motivate us—force us, I think—not to identify properties with the sets of their instances. If we identify them instead with functions from worlds to extensions, we get a theory of properties that is neutral with respect to disputes over counterpart theory, and we avoid a problem for Lewis's theory of events.

Similar problems about *temporary* second-order predication motivate us—though this time I don't think that they *force* us—to give up this theory as well, and to identify properties with functions from $\langle \text{world}, \text{time} \rangle$ pairs to extensions. Again, the replacement theory is neutral with respect to a metaphysical dispute that the old theory (arguably) forces us to take a stand on—the dispute over whether objects have temporal parts. It also allows us to give a smoother semantics for predication, to better accommodate our intuitions about which objects temporary properties are properties of, and to make temporally self-locating beliefs genuinely *self-locating*.

Australian National University

REFERENCES

- Balashov, Yuri 1999. Relativistic Objects, *Noûs* 33: 644–62.
 Cartwright, Richard 1975. Scattered Objects, in *Analysis and Metaphysics*, ed. Keith Lehrer, Dordrecht: Reidel: 152–71.
 Chomsky, Noam 2000. *New Horizons in the Study of Language and Mind*, Cambridge: Cambridge University Press.

³² In fact, people probably don't have many beliefs at all, if we adopt a temporal parts based solution to the problem of change. Properties like *believing that P* will almost always be temporary, and so they won't be properties of continuant people, but of their stages. There is, of course, footwork to be done here—there are ways of making it out that, even though people never instantiate belief-properties, people still *have beliefs* by having parts that instantiate belief-properties. We can do this, but it would be nice if we didn't have to.

³³ Again, it's a bit tricky to say carefully what's going on. Just as we can treat functions from worlds to extensions as sets of ordered pairs, we can treat a function f from $\langle \text{world}, \text{time} \rangle$ pairs to extensions as the set of $\langle w, t, i \rangle$ triples (where w is a world, t is a time, and i is an object) such that i is a member of $f(\langle w, t \rangle)$. Then we can think of your self-locating beliefs (including your temporally self-locating beliefs) as attitudes toward sets of such triples, such that the belief is accurate iff $\langle @, \text{now}, \text{you} \rangle$ is in the set. This mirrors accounts of centred-worlds contents in which centred worlds include a specification of the time. (Thanks to Juan Comesaña for discussion here.)

- Haslanger, Sally 1989a. Persistence, Change, and Explanation, *Philosophical Studies* 56: 1–28.
- Haslanger, Sally 1989b. Endurance and Temporary Intrinsic, *Analysis* 49: 119–25.
- Heller, Mark 1998. Property Counterparts and Ersatz Worlds, *Journal of Philosophy* 95: 293–316.
- Lewis, David 1968. Counterpart Theory and Quantified Modal Logic, *Journal of Philosophy* 65: 113–26.
- Lewis, David 1971. Counterparts of Persons and their Bodies, *Journal of Philosophy* 68: 203–11.
- Lewis, David 1973. *Counterfactuals*, Oxford: Blackwell.
- Lewis, David 1979. Attitudes De Dicto and De Se, *Philosophical Review* 88: 513–43.
- Lewis, David 1983a (1970). General Semantics, in *Philosophical Papers 1*, New York: Oxford University Press: 189–232.
- Lewis, David 1983b. Postscript B to ‘Survival and Identity’, in *Philosophical Papers 1*, New York: Oxford University Press: 73–7.
- Lewis, David 1986a. *On the Plurality of Worlds*, Oxford: Blackwell.
- Lewis, David 1986b. Events, in *Philosophical Papers 2*, Oxford: Oxford University Press: 241–69.
- Lewis, David 1990. *Parts of Classes*, Oxford: Blackwell.
- Lewis, David 1999 (1983). New Work for a Theory of Universals, in *Papers in Metaphysics and Epistemology*, Cambridge: Cambridge University Press: 8–55.
- Lewis, David 2002. Tensing the Copula, *Mind* 111: 1–13.
- Sider, Theodore 2001. *Four Dimensionalism*, Oxford: Clarendon Press.
- Yablo, Stephen 1998. Intrinsicness, *Philosophical Topics* 26: 479–504.