

Almost Identical, Almost Innocent

Katherine Hawley

Abstract

In his 1991 book, *Parts of Classes*, David Lewis discusses the idea that composition is identity, alongside the idea that mereological overlap is a form of partial identity. But this notion of partial identity does nothing to help Lewis achieve his goals in that book. So why does he mention it? I explore and resolve this puzzle, by comparing *Parts of Classes* with Lewis's invocation of partial identity in his 1993 paper 'Many But Almost One', where he uses it to address Unger's problem of the many. I raise some concerns about this way of thinking of partial identity, but conclude that, for Lewis, it is an important defence against accusations of ontological profligacy.

1. Composition as Identity

David Lewis's *Parts of Classes*, published in 1991, is a sustained argument that we should understand set theory in terms of mereology, the theory of parts and wholes. 'The notion of a singleton, or unit set, can serve as the distinctive primitive of set theory. The rest is mereology: a class is the fusion of its singleton subclasses, something is a member of a class iff its singleton is part of that class'.¹ The technical aspects of this project travelled a somewhat rocky road. The book itself contains a last-minute appendix co-written with John P. Burgess and A.P. Hazen, whilst just two years later, Lewis published a follow-up paper, which he described as 'an abridgement of parts of *Parts of Classes* not as it is, but as it would have been had I known sooner what I know now'.²

Parts of Classes is currently out of print. But one particular non-technical section of the book has become a break-out hit, much discussed as a potential contribution to our understanding of mereology itself, as opposed to our understanding of set theory in

¹ David Lewis, *Parts of Classes* (Oxford, and Cambridge MA: Basil Blackwell Ltd., 1991) vii.

² David Lewis, 'Mathematics is Mereology', *Philosophia Mathematica*, 1, 3–23, at 3. John P. Burgess, 'Lewis on Mereology and Set Theory', in *A Companion to David Lewis*, edited by Barry Loewer and Jonathan Schaffer (Malden MA: Wiley-Blackwell, 2015), is a user-friendly guide to Lewis's writing on these topics, and includes a fetching attempt to make 'mereology' happen. Karen Bennett, "'Perfectly Understood, Unproblematic, and Certain" Lewis on Mereology', in the same volume, explores many of the themes I touch on in the present article.

terms of mereology. This is section 3.6, the few pages entitled ‘Composition as Identity’, in which Lewis argues that mereology is ontologically innocent.

In this context, ‘mereology’ means a specific, demanding theory of parts and wholes, one which includes the axiom of Unrestricted Composition: whenever there are some things, then there exists a fusion of those things. (If we take the notion of part as basic, then by definition, something is a fusion of some things iff it has all of them as parts and has no part that is distinct from each of them.)³ The axiom of Unrestricted Composition looks far from innocent – it appears to be the very model of a worldly claim. After all, it posits an additional entity (a fusion) corresponding to every plurality of entities. And these worldly consequences are not incidental to Lewis’s invocation of that axiom: he needs the multitudinous entities in order to make sure that the universe is furnished generously enough for standard set theory.

Lewis needs mereology to provision us with a host of fusions, but he realises that this multitudinousness is likely to raise suspicions. One kind of suspicion focuses on the sheer weirdness of the arbitrary fusions to which we are committed by Unrestricted Composition. There are, for example, the trout-turkeys, each of which is composed of an undetached half-trout and a spatially distant undetached half-turkey. But Lewis addresses doubts about trout-turkeys in the preceding section of his chapter, and section 3.6 is addressed not to the weirdness of arbitrary fusions, but to their sheer number. Surely it is a huge strike against the theory of mereology that it is so ontologically profligate?

Lewis argues that, despite appearances, mereology is ontologically innocent, not profligate. On what grounds? The most (in)famous idea in this section is that the relation which holds between the cats, for example, and the fusion of the cats, is the relation of identity; well, sort of. Lewis writes ‘The fusion is nothing over and above the cats that compose it. It just *is* them. They just *are* it. Take them together or take them separately, the cats are the same portion of Reality either way’.⁴ As is standard, let’s call the relation which holds between the many cats and the fusion of those cats

³ Lewis, *Parts of Classes*, 73–74.

⁴ Lewis, *Parts of Classes*, 81.

‘composition’. Lewis says he is endorsing what he calls the ‘Thesis of Composition as Identity’.

2. Partial Identity

A recent collection of essays, also entitled *Composition as Identity*, explores the metaphysical and logical challenges which confront anyone hoping to identify the many cats with the single cat-fusion.⁵ *Parts of Classes* is cited in almost all of the essays, and my own contribution to that volume focuses closely on Lewis’s remarks about ontological innocence and the thesis of Composition as Identity.⁶ But I, along with the other contributors, neglected one particular puzzling aspect of Lewis’s discussion in his section 3.6.

After arguing that ‘[t]he “are” of composition is, so to speak, the plural form of the “is” of identity’, Lewis writes ‘[i]n endorsing Composition as Identity, I am following the lead of D.M. Armstrong and Donald Baxter’.⁷ Baxter’s work is clearly relevant to Lewis’s discussion: it is the main contemporary source for the claim that composition is identity, a claim which, if vindicated, would certainly ground the ontological innocence of mereology.⁸ If a theory commits you only to entities which are identical to entities you already accept, then it brings you no additional commitments: it is ontologically innocent.

But what about Armstrong? As Lewis says, ‘Armstrong takes strict identity [i.e. identity] and strict difference [i.e. complete disjointness] as the endpoints of a spectrum of cases, with cases of more or less extensive overlap in between.’⁹ Lewis quotes Armstrong on the relationship between two terraced houses which share a common party wall:

[they] are not identical, but they are not completely distinct from each other either. They are partially identical, and this partial identity takes the form of having a common part. Australia and New South Wales are not identical, but

⁵ A.J. Cotnoir and Donald L.M. Baxter (eds.) *Composition as Identity* (Oxford: Oxford University Press 2014).

⁶ Katherine Hawley, ‘Ontological Innocence’, in *Composition as Identity*, 70–89.

⁷ Lewis, *Parts of Classes*, 82.

⁸ Donald L.M. Baxter, ‘Identity in the Loose and Popular Sense’, *Mind* 97 (1988), 375–82, and Donald L.M. Baxter, ‘Many-One Identity’, *Philosophical Papers* 17 (1988), 193–216.

⁹ Lewis, *Parts of Classes*, 82,

they are not completely distinct from each other. They are partially identical, and this partial identity takes the form of the whole-part ‘relation’...¹⁰

Partial identity in this sense is a matter of incomplete mereological overlap. According to Lewis’s (standard) definition, objects overlap iff they have some common part.¹¹ If objects have all their parts in common, then they completely overlap and so they are literally, strictly, completely identical, at least according to the extensional mereology Lewis favours. If objects have no parts in common, then they are not at all identical; they are ‘disjoint’. And the idea is that, on a sliding scale between these two extremes, objects which share some but not all of their parts are partially identical, identical to some non-zero degree; the more parts in common, the more identical they are.

Why does Lewis mention Armstrong’s partial identity in the same breath as Baxter’s composition as identity? Overlap is not composition. Two terraced houses do not compose one another, and more generally overlap is symmetric but composition is not. Composition typically holds between several things and one thing; only in the limit case of self-composition does it hold one-one. Composition is a multi-grade relation, although this point is harder to formulate if composition literally is identity. Overlap is defined as holding between one object and one object. Several objects might each stand in the overlap relation to a single thing, but this does not entail that they plurally somehow overlap it, although their fusion, if it exists, overlaps it.¹²

What about the special case of overlap between an object and one of its proper parts, as between Australia and New South Wales? In such a case, every part of the smaller object is a part of the larger one, but not vice versa. Even here, the overlap relation is not the same as the composition relation. Overlap holds between each of the parts and the larger object (between New South Wales and Australia, between Queensland and Australia, and so on). It does not hold between the various parts collectively (the states and territories of Australia) and the larger object (Australia).

¹⁰ D.M. Armstrong, *Universals and Scientific Realism II* (Cambridge: Cambridge University Press, 1980), 37–38.

¹¹ Lewis, *Parts of Classes*, 73.

¹² Peter van Inwagen, ‘Composition as Identity’, *Philosophical Perspectives* 8 (1994), 207–220 has illuminating remarks about Lewis’s terminology here, in his footnote 2.

So Armstrong's relation of partial identity is importantly different from the composition relation which, according to Baxter, is the identity relation, even though Lewis apparently lumps them together. It seems that only Peter van Inwagen has noted this oddity in print. In his 1994 paper 'Composition as Identity', van Inwagen attacks Lewis's claim that the thesis of unrestricted composition is ontologically innocent. He launches a broadside against the Baxter-style thesis that composition is identity, i.e. that the whole just is the parts, charging it with incomprehensibility. But van Inwagen is more positive when he considers Lewis's citation of Armstrong's partial identity:

...there does seem to be some sort of insight to be gleaned from thinking of overlap as partial identity...calling overlap 'partial identity' is a perspicuous way of calling attention to the fact that the roots of both parthood and identity are contained in the overlap relation.¹³

But as he points out 'someone could accept Armstrong's thesis about the intimate relation between overlap and identity without being in any way moved to regard Mereology as ontologically innocent'.¹⁴ Imagine, says van Inwagen, that you accept the existence of various objects, the *as*, but deny that they collectively compose anything: you deny unrestricted composition. A Lewis-style philosopher tries to persuade you that accepting the existence of *A*, the sum of the *as*, would in some sense be ontologically innocent, no genuine addition to the ontology you already accept.

A Baxter-style approach, if we could make sense of it, would vindicate the idea of ontological innocence: *A* just is the things you already accept, the *as*, no addition to your catalogue of the universe. But how could partial identity help here? If *A* exists, then each of the *as* is partially identical to *A*. But none of the *as* is fully identical to *A*, even if *A* does exist. So accepting *A* into your ontology expands that ontology, by adding an object which is not fully identical to anything you already accept. There may be good reason to accept this expansion, but expansion it is nevertheless.

¹³ Van Inwagen, 'Composition as Identity', 214–5.

¹⁴ Van Inwagen, 'Composition as Identity', 216.

Following Armstrong by regarding overlap as partial identity doesn't help Lewis in his project of showing that the thesis of unrestricted composition is ontologically innocent. But this assimilation of overlap to identity does important work for Lewis elsewhere, in ways which relate to his concern with ontological innocence, parsimony, and profligacy.

3. Many But Almost One

The most explicit discussion is in 'Many, But Almost One', which was first published as part of a *festschrift* for Armstrong in 1993, i.e. not long after *Parts of Classes*.¹⁵ This paper focuses on the problem of the many, which arises in connection with the seeming fuzziness of ordinary objects.

As we usually think of them, individual cats do not have sharp boundaries. They are constantly ingesting, digesting, and excreting, breathing in and out, and shedding hairs. If we insist that there is no vagueness in the world, so that objects must have sharp boundaries, yet we do not want to deny the existence of cats, then we face various problems. One concern is that fuzziness may seem built into the very concept of cattiness, so that a sharply-bounded cat is a conceptual impossibility. However I will set this issue aside to focus on the problem which concerned Lewis, and before him Peter Unger.¹⁶

As Lewis points out, there are lots of different aggregates of particles which have a decent claim to make up the cat. If we allow Lewis to slide from talk of particles plurally, to talk of aggregates, to talk of composite objects, then it looks as if there are very many precise, overlapping, cat-like objects curled up on the mat, each differing from the others in only a few small parts around its edges. All are equally good candidates for being the cat. So which is the cat?

Lewis first develops a supervenient story: there is no single candidate of which we can truly say that it is a cat, but nevertheless it is true to say that exactly one of the candidates is a cat. 'Cat' is an imprecise concept, but for each permissible way of

¹⁵ David Lewis, 'Many, But Almost One', pages references as republished in his *Papers on Metaphysics and Epistemology* (Cambridge: Cambridge University Press 1999).

¹⁶ Peter Unger, 'The Problem of the Many', *Midwest Studies in Philosophy*, 5 (1980): 411–67.

sharpening up this concept, it will pick out exactly one of the candidates to the exclusion of all the others. This approach is controversial, but we can see how, if it is successful, it allows us to assert that there is exactly one cat on the mat (since this is true no matter how we sharpen up ‘cat’), and yet resist the idea that the world itself is vague or indeterminate.

Lewis then wheels in partial identity, to provide a second, different solution to the problem of the many. He here cites the same Armstrong passage that he quotes in *Parts of Classes*.¹⁷ In his brief response to Lewis’s ‘Many, But Almost One’, Armstrong seems to welcome this intervention, and he mentions being influenced by F.H. Bradley’s talk of partial identity.¹⁸ As I understand matters, Bradley was engaged with William James about what it is for sense impressions to resemble one another. The proposal is that partially resembling sense impressions literally share some parts, and thus can be considered partially identical. Following this tradition, discussions of partial identity in present-day philosophical literature usually focus on instantiation relations between universals and particulars, or else relations amongst universals.¹⁹ In this paper, however, I focus upon partial identity amongst material particulars, not universals.

Lewis suggests that each of the cat-candidates is partially identical with each of the others. Moreover, this partial identity is much closer to the identity end of the spectrum than to the disjointness end, since the cat-candidates overlap one another almost completely, sharing almost all of their parts. Lewis writes:

Assume our cat-candidates are genuine cats. (Set aside, for now, the supervaluationist solution.) Then, strictly speaking, the cats are many. No two of them are completely identical. But any two of them are almost completely

¹⁷ Lewis, ‘Many But Almost One’, 176.

¹⁸ D.M. Armstrong, ‘Reply to Lewis’, in *Ontology, Causality and Mind: Essays in Honour of D.M. Armstrong*, edited by John Bacon, Keith Campbell and Lloyd Reinhardt (Cambridge University Press, 1993): 38–42. F.H. Bradley ‘On Professor James’ Doctrine of Simple Resemblance’, *Mind* 2 (1893), 83–88.

¹⁹ For examples, see Arada Denkel, ‘Resemblance Cannot Be Partial Identity’, *Philosophical Quarterly* 48 (1998): 200–204, D.L.M. Baxter, ‘Instantiation as Partial Identity’, *Australasian Journal of Philosophy* 79.4 (2001): 449–64, and Nicholas Mantegani, ‘Instantiation is Not Partial Identity’, *Philosophical Studies* 163.3 (2013): 697–715.

identical; their differences are negligible, as I said before. We have many cats, each one almost identical to all the rest...²⁰

How does this help solve the problem of the many? The idea is that it legitimates our intuitive judgement that there is exactly one cat on the mat, by somehow showing that judgement to be almost true, or otherwise acceptable. ‘The cats are many, but almost one. By blameless approximation, we may say simply that there is one cat on the mat. Is that true? – Sometimes we’ll insist on stricter standards, sometimes we’ll be ambivalent, but for most contexts it’s true enough.’²¹

Lewis is gesturing towards a kind of contextualism here. His influential ‘Elusive Knowledge’ was first published three years later than ‘Many But Almost One’.²² There, he presents a more carefully worked-out form of contextualism about knowledge ascriptions.²³ So what might be the contextual element to ‘there is one cat on the mat’? As Lewis explains, it’s not the predicate ‘is a cat’ which is subject to contextually varying standards – there’s no particular object which counts as a cat when we are being casual, but not when we’re being strict. Instead, it is the underlying identity claim which is somehow subject to contextual variation.

Remember how we translate statements of number into the language of identity and quantification. ‘There is one cat on the mat’ becomes ‘For some x , x is a cat on the mat and every cat on the mat is identical to x ’. That’s false, if we take ‘identical’ to express the complete and strict identity that lies at the end of the spectrum. But the very extensive overlap of the cats does approximate to complete identity. So what’s true is that for some x , x is a cat on the mat, and every cat on the mat is almost identical to x . In this way, the statement that there is one cat on the mat is almost true.²⁴

There are at least two ways of developing this kind of picture. One thought is that ‘there is one cat on the mat’ is false in every context, but that in many contexts it’s

²⁰ Lewis, ‘Many But Almost One’, 178.

²¹ Lewis, ‘Many But Almost One’, 178.

²² Lewis, ‘Elusive Knowledge’, *Australasian Journal of Philosophy*, 74, 549–67.

²³ Jonathan Schaffer, ‘Lewis on Knowledge Ascriptions’, in Loewer and Schaffer (eds.) *A Companion to David Lewis* (Wiley-Blackwell 2015) is a helpful guide.

²⁴ Lewis, ‘Many But Almost One’, 178.

permissible to assert this false claim: assertion conditions vary with context. A different thought is that ‘there is one cat on the mat’ is false in strict contexts, but that nevertheless in many contexts it is true: truth conditions vary with context. There are surely important differences between these two pictures, and on either approach we’d need to do more work if we hoped to capture Lewis’s talk of ‘true enough’ and ‘almost true’. But I will not dwell on the details here, since my main goal is to understand how the notion of overlap as partial identity is supposed to license some or other contextual variation of this kind.

So: how does talk of the partial identity of the various cat-candidates help Lewis, over and above talk of the mereological overlap of the cat-candidates? Presumably the idea is that the closer things are to being absolutely, completely identical, the easier it is for there to be contexts in which it is acceptable simply to say that they are identical, that there is just one of them. Consider Mog and Bagpuss, two cats which have no parts in common, and live in different cities. There is no context in which it is appropriate to say that Mog and Bagpuss are one and the same cat, that there is just one of them. Now consider C’ and C’’, two precisely-bounded cat candidates which share almost all of their parts in common. In a context – perhaps a seminar discussion of metaphysics – in which we are governed by very strict standards, we should not say that C’ and C’’ are one. But in most other contexts, we can permissibly say that they are one.

This line of thought suggests that ‘partial identity’ is not just a quirky term for ‘mereological overlap’. Instead the supposed scale from disjointness, on through increasing overlap, and up to identity is to be understood as genuinely identity-like, in that it supports the attribution of identity in increasingly stricter contexts, culminating in the strictest context with what we might think of as real identity. Partial identity is used as the basis of a notion of almost-identity which is then used to account for, or approximately preserve, our intuitive judgements of identity.

There are reasons to hesitate before following Lewis (and Armstrong) down this road. When we picture the sliding scale from disjointness to full identity, it is tempting to picture something like a sequence of Venn diagrams, beginning with two circles which have no overlap, then increasingly large intersections, and finally one circle in

the middle of the page. That's to say, it is tempting to picture overlap in spatial terms. On this picture, it is easy to accept that there is indeed a sensible sliding scale between complete disjointness to identity, that we can order various pairs of objects by the degree of overlap between them. Some pairs will be equal in their degree of overlap, but none will be incommensurable in this respect.

But what is the relationship between mereological overlap and spatial overlap?

Objects which share some spatial parts overlap spatially, so long as we ignore peculiar issues about time-travel or multilocation.²⁵ And objects which overlap spatially share some spatial parts, so long as we ignore the alleged possibility of coincidence without such sharing of parts, as Lewis usually does. Even allowing for these ignorings, a further assumption is required if we are to align degree of spatial overlap with degree of mereological overlap. The assumption is that these march hand-in-hand, because each further increment of mereological overlap increases spatial overlap by the same degree. To put it bluntly, we need to assume that all of an object's smallest parts are the same size as one another.

To see what happens if we drop this assumption, imagine an object Orlando which has both small and large objects amongst its most basic parts, i.e. amongst those parts which do not have parts of their own. Orlando shares all its basic parts except for one large one with a second object, Orla. Mereologically, Orlando and Orla are almost entirely 'identical', but there is a large difference between their spatial extents, because the one part in which they differ is a big one. Orlando also shares all of its basic parts except for two tiny ones with a third object, Otis. Mereologically, Orlando and Otis overlap less than do Orlando and Orla – they differ in twice as many parts – but Orlando and Otis occupy almost exactly the same spatial region as one another, because the parts in which they differ are so small.

Which pair are nearer to being just one and the same thing? If mereological overlap is partial identity, then Orlando and Orla are nearly identical, more so than Orlando and Otis. But Orlando and Otis will *look* more nearly identical, because they match so

²⁵ Useful surveys of this material include Cody Gilmore, 'Location and Mereology', *Stanford Encyclopedia of Philosophy* ed. E. Zalta (Fall 2017) <https://plato.stanford.edu/archives/fall2017/entries/location-mereology/> and Shieva Kleinschmidt 'Introduction' to *Mereology and Location*, ed. Kleinschmidt (Oxford University Press, 2014).

closely in their spatial extents. This isn't *per se* a problem for the view that mereological overlap is partial identity. But it is a warning not to be lulled into accepting that view too easily, not to slip unnoticed from intuitions about spatial overlap to intuitions about mereological overlap.²⁶

In fact, this point about the difference between spatial and mereological overlap, and their supposed relations to partial identity, is a special case of something Lewis himself draws to our attention. Real, strict identity is conceptually tied to indiscernibility: it may be controversial whether indiscernibility guarantees identity, but it is uncontroversial that identity guarantees indiscernibility. Yet partial identity, even almost-identity, does not stand in any straightforward relationship to indiscernibility.

Lewis writes:

Further, the relation of almost-identity, closely akin to the complete identity that we call identity *simpliciter*, is not a relation of partial indiscernibility. Of course we can expect almost-identical things to be very similar in a great many ways: size, shape, location, weight, purring, behaviour, not to mention relational properties like location and ownership. But it is hard to think of any very salient respect in which almost-identical things are guaranteed to be entirely indiscernible.²⁷

Notice that if degree of mereological overlap can come apart from degree of spatial overlap, then degree of spatial overlap may be the better guide to similarity of ordinary properties. Moreover, depending upon our other metaphysical commitments, we may think that there are also much more significant differences between objects which differ only marginally in their parts. Sortal properties such as *being a cat*, or *being a car*, are often regarded as 'maximal', which is to say that a condition for instantiating the property is that the object in question not be a large part of something else which instantiates that property. The object (if there is one) which is my car without its right wing mirror would itself be a car, if it were not attached to the wing

²⁶ There may also be more complicated reasons to question the possibility of ordering by degree of mereological overlap, if gunk is actual; see Hud Hudson, 'Simples and Gunk', *Philosophy Compass* 2.2 (2007) 291–302.

²⁷ Lewis, 'Many But Almost One', pp. 178–9.

mirror; instead, it is a large proper part of a car. On this view, small differences in composition can make a very significant difference in sortal property.

So the sliding scale picture is not as straightforward as it first seems, and there is no straightforward relationship between partial identity and partial indiscernibility, even though the relationship between identity and indiscernibility is so intimate. Even if we set these worries aside, it is not clear what Lewis gains by referring to mereological overlap as ‘partial identity’. The sketch of contextualism doesn’t seem to require that mereological overlap be any kind of identity, so long as there is an underlying structure which can support contextual variation in the assertability of ‘there is one cat on the mat’.

4. Parsimony and Profligacy

Why then is Lewis so keen on regarding mereological overlap as a weakened or partial form of identity? I think the best way of understanding this aspect of both *Parts of Classes* and ‘Many But Almost One’ is in terms of Lewis’s concern with ontological parsimony as a criterion for theory choice. The choice in question is between a doctrine of Unrestricted Composition, and a more moderate, commonsensical account of composition according to which some, but not all, pluralities have fusions. (A nihilist third option, denying the existence of absolutely all fusions, was not a credible rival when Lewis was writing.)

This suggestion may seem puzzling. Following van Inwagen, I have already argued that the equation of mereological overlap with partial identity does not help establish the ontological innocence of the doctrine of Unrestricted Composition. None of the many cats is almost identical to the fusion of all cats, and that’s what would be needed to make Lewis’s point in the context of *Parts of Classes*. But there are somewhat different questions of parsimony and profligacy at stake in Lewis’s treatment of the problem of the many.

Setting aside the supervenient option, the problem of the many is supposed to be that we seem to have too many cats on the mat(s). Recognising overlap as partial identity is supposed to resolve this by allowing that the many cats are almost just one and the same thing, in such a way that a dose of contextualism often permits us to say

simply that they are the same thing. This formulation of the problem of the many has the doctrine of Unrestricted Composition as a background assumption. After all, what other reason is there to believe in all the catlike cat-candidates in the first place? Admittedly, there is a version of the problem of the many which focuses on pluralities of particles, pointing out that there seems to be no unique such plurality which makes up ‘the’ cat at any given moment. But in this version, the problem has an attractive solution: there is one cat, which has indeterminate boundaries because it is indeterminate exactly which things are its parts.²⁸

Lewis famously rejects the idea that the world itself may be vague or indeterminate, and explicitly discusses this in connection with composition.²⁹ There seems to be no sensible way of drawing a sharp boundary between cases of composition, and cases where composition does not occur, and so he opts for universalism: every plurality has a fusion. Thus he is committed to the more problematic version of the problem of the many – the one which features many competing cats, not just many competing pluralities of cat-parts – because of his rejection of ontic indeterminacy, and his consequent acceptance of the doctrine of Unrestricted Composition. This in turn generates an important dialectical role for the idea that partial overlap amounts to partial identity.

Lewis is committed to very many more catlike objects than is his opponent. But if these ‘additional’ objects are almost identical to the object (the cat) his opponent does accept, then they barely add anything to the ontological cost of the theory. Suppose that we accept the connections between overlap, almost-identity, and thereby to ontological almost-innocence. What significance does this have for metaphysical debate about what’s on the mat?

The connections do not provide positive reason for accepting the doctrine of Unrestricted Composition. To see this, imagine that you have accepted just one of the catlike things into your ontology. This might be because you have a ‘brutal’ view of composition, thinking that there is a brute, unknown fact of the matter as to which

²⁸ This type of ontic indeterminacy does not lead inevitably to ontic indeterminacy in existence, as I discuss in ‘Vague Existence’, *Proceedings of the Aristotelian Society* CII (2001-02), 125–40.

²⁹ David Lewis, *On the Plurality of Worlds* (Basil Blackwell 1986), 212–13.

single collection of cat-parts composes a cat-sized object.³⁰ Relatedly, an epistemicist approach to indeterminacy in composition could bring you here, although epistemicism about indeterminacy in ‘is a cat’ won’t do the job alone.³¹

Now, I ask you to accept another such precise, catlike thing into your ontology, and to sweeten the pill I tell you that this new thing is almost identical to the one you already accept. The addition is almost no addition at all; accepting it is almost an ontologically innocent move. After all, I’m not asking you to accept a whole new distinct object into your ontology, merely one which massively overlaps with something you already accept. If I tell you that this move is ontologically innocent, then perhaps what I say is approximately or almost true. But even if you accept that this move would be almost innocent, that doesn’t positively recommend it; not making the move would be entirely innocent, which is surely preferable. And if you have a brutalist view of composition, you will struggle to reconcile it with this two-but-almost-one cats picture. Moreover, I won’t stop at two cats: I will try to persuade you that addition after addition is acceptable since almost innocent. Small additions will quickly sum to large, not-so-innocent additions.

Connecting almost-identity to ontological innocence does not give us positive reasons to accept that there are many cat-like objects on the mat(s). Rather, almost-identity is a defensive tool available to theorists – like Lewis – who claim that there are very many cat-like objects on the mat. If all these objects are almost-identical to one another, and if almost-identity really implies almost-innocence, then the theory is almost not profligate.

This connection between mereological overlap and ontological innocence also does significant work for Lewis elsewhere.³² In brief, Lewis’s account of persistence combines a temporal-parts ontology with the doctrine of Unrestricted Composition, populating the world with an enormous number of overlapping, criss-crossing, diverging objects. Indeed, it is this huge population which gives the theory the flexibility to account for indeterminacy in persistence, for fission and fusion, and so

³⁰ Ned Markosian, ‘Brutal Composition’, *Philosophical Studies* 92.3 (1998), 211–249.

³¹ Timothy Williamson, *Vagueness*, (London: Routledge, 1994).

³² Katherine Hawley, ‘David Lewis on Persistence’, in *A Companion to David Lewis*, Loewer and Schaffer (eds.) (2015).

on; this flexibility is touted as a significant reason to accept perdurantism in this form. However, in principle the same flexibility can be obtained without temporal parts, if we are willing to accept a plenitudinous endurance theory, which populates the world with a huge number of coincident objects, each persisting without being temporal parts of one another.³³ But if mereological overlap reduces ontological commitment, then the plenitudinous perdurantist has a key advantage over the plenitudinous endurantist, one which needs to be weighed in the balance of theory choice. After all, the supposed perduring objects are almost identical with one another, unlike the supposed enduring objects.

5. Conclusion

In this paper I set out to explore a quirk in Lewis's famous discussion of composition as identity. He integrates discussion of the sense in which a composite object 'just is' its parts, with discussion of the sense in which two mostly-overlapping objects are 'almost identical'. But these are quite distinct claims, which makes it *prima facie* puzzling why they are discussed together. I explored connections with Lewis's use of almost-identity in addressing the problem of the many, and raised some concerns about the sliding-scale picture which seems to underlie the assimilation of mereological overlap to identity. But this assimilation is most powerful when seen as a defensive move for advocates of unrestricted composition: it lessens the ontological burdens of that view, as compared to a supposedly more commonsensical moderate account of composition. Although partial identity doesn't help Lewis establish the ontological innocence of many-one composition in the context of Parts of Classes, it nevertheless speaks to his wider concerns about ontological commitment, parsimony, and profligacy.

Acknowledgements: I am grateful to Aaron Cotnoir for his helpful advice.

University of St Andrews
kjh5@st-andrews.ac.uk

³³ As explored by e.g. John Hawthorne, *Metaphysical Essays* (Oxford: Oxford University Press 2006), and Ross Inman, 'Neo-Aristotelian Plenitude', *Philosophical Studies* 168.3 (2014), 583–97.