

Ontological Innocence¹

Katherine Hawley, March 2013. For *Composition as Identity*, edited by Donald Baxter and Aaron Cotnoir, forthcoming with Oxford University Press.

Abstract

In *Parts of Classes*, David Lewis argues that mereology is ontologically innocent, and connects this to the thesis that composition is identity. This chapter investigates whether mereology can be regarded as ontologically innocent even if composition is not identity. One idea is that we are all implicitly committed to the existence of arbitrary sums even before we accept mereology, so that accepting mereology does not give us any new commitments. A different idea is that, although accepting mereology gives us new ontological commitments, the additional entities do not add explanatory complexity, so do not offend against considerations of parsimony.

Keywords: mereology; composition; ontological innocence; ontological commitment; parsimony; David Lewis.

Read it again:

...Mereology is ontologically innocent.

To be sure, if we accept mereology, we are committed to the existence of all manner of mereological fusions. But given a prior commitment to cats, say, a commitment to cat-fusions is not a *further* commitment. 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. Commit yourself to their existence all together or one at a time, it's the same commitment either way. If you draw up an inventory of Reality according to your scheme of things, it would be double counting to list the cats and then also list their fusion. In general, if you are already committed to some things, you incur no further commitment when you affirm the existence

¹ I presented this material at Rutgers Metaphysical Mayhem 2012, and am very grateful for the support I received there whilst working through my issues. Thanks in particular to Einar Bohn, Donald Baxter, Aaron Cotnoir, Jonathan Schaffer, and Ted Sider. This work was supported by the European Community's Seventh Framework Programme FP7/2007-2013, under grant agreement no. FP7-238128.

of their fusion. The new commitment is redundant, given the old one. (Lewis 1991, 81-82)

There is something very attractive about much of this. And yet for many of us, perhaps even for Lewis, there is something rather repulsive about the thesis that composition is identity. Can we salvage the attractive elements, whilst rejecting the repulsive? My goal is to extricate and examine Lewis's ideas about ontological innocence, commitment and double counting, attempting to understand these primarily as epistemic or methodological claims. How far can we get down this route without adopting radical metaphysical theses about composition as identity?

1. Ontological Innocence

Lewis's main purpose in this section of *Parts of Classes* is to show that mereology is ontologically innocent. In this context, 'mereology' means the thesis that composition is unrestricted; this contrasts with moderate views of composition, according to which composition occurs in some cases but not in others, and with nihilist views, according to which composition never occurs. What does it mean to say that mereology is ontologically innocent? A clue: in this respect, mereology is like plural quantification and unlike set theory.

When we quantify plurally over ordinary things, we do not, according to Lewis, secretly quantify over sets or classes of those things. When we say that the fans of the Chieftains are many, we are not saying, of the set which has all and only the Chieftains fans as members, that it has the property of being many-membered. Instead, we are talking about the individual fans of the Chieftains, and saying that they are many.

[Plural quantification] is not ordinary singular quantification over special plural things...Plural quantification, like singular, carries ontological commitment only to whatever may be quantified over. It is devoid of set theory and it is ontologically innocent. (Lewis 1991, 68-9, see also 102).

If you are already committed to the existence of some objects, perhaps by accepting sentences which quantify over them in a singular way, then you do not take on any extra ontological commitments when you accept sentences which quantify over those same things again in a plural way. After all, you are just quantifying over those same

things again. (As is standard, I will take it that ontological commitment attaches primarily to theories or sentences, and derivatively to the people who believe such theories or sentences.)

Set theory, in contrast, is not ontologically innocent, because of singleton-formation.

Set theory is not innocent. Its trouble has nothing to do with gathering many into one. Instead, its trouble is that when we have one thing, then somehow we have another, wholly distinct thing, the singleton. And another, and another...*ad infinitum*. But that's the price for mathematical power. Pay it. (Lewis 1991, 87)

If you are already committed to the existence of some objects, and you then accept set theory, you thereby take on extra ontological commitments, including commitment to the existence of the various singletons of the various objects you were already committed to. (And so on *ad infinitum*.) The singletons were not previously in your ontology, but now they are: this shift, induced by accepting set theory, is thus not ontologically innocent.

So ontological innocence is a kind of ontological conservativeness: you can accept an ontologically innocent claim without thereby adding to the list of entities whose existence you are committed to. This suggests that whether a theory or sentence is ontologically innocent may be a relative matter: whether it carries *additional* commitments for you may depend upon what ontological commitments you already have. But this relativity plays no role in the case of mereology. The thesis of unrestricted composition does not unconditionally specify what composite objects there are; instead, it makes conditional claims e.g. that if there are some cats, then there is the fusion of those cats. Mereology is ontologically innocent if and only if the claim that there is a fusion of (e.g.) the cats is ontologically innocent for anyone who has already accepted that the cats exist. Likewise, moderate views of composition involve conditional existence claims which may or may not be ontologically innocent.

Why does it matter whether mereology is ontologically innocent? For Lewis, the ontological non-innocence of set theory is a cost, albeit one worth bearing in order to receive the benefits of mathematical power. This suggests a methodology of theory choice based on criteria including ontological parsimony, in line with Lewis's

methods and methodological reflections elsewhere (e.g. Lewis and Lewis (1970); for discussion see Nolan (2005, chapter 9) and Nolan (forthcoming, section 3)). When choosing between theories, we should weigh up the costs and benefits of each option, for example the relative unity and economy of each theory. Other things being equal, a theory which involves fewer ontological commitments is preferable to one which involves more. If mereology is ontologically innocent, then accepting mereology does not add to our ontological commitments; this removes one obstacle to accepting mereology.

But why think that mereology is ontologically innocent? If composition is identity, then ontological innocence is secured. Suppose that you are committed to the existence of some objects; you are thereby committed to the existence of any object(s) they are identical to. If composition is identity, then you can now accept a theory according to which those objects compose something, without thereby expanding your ontology. When you tally up your ontological commitments, having accepted the existence of the composite object, it would be a mistake first to count the original objects, and then to increase the total by adding an object – the composite – which is identical to those you have already counted. That would be double-counting: a miser does not increase his wealth by counting the same pennies over and over again.

If composition is identity, then any claim that some objects have a sum is ontologically innocent, given a prior commitment to those initial objects. So if composition is identity, then the thesis of unrestricted composition is ontologically innocent, as are more moderate claims that certain pluralities of objects have sums, though others do not. Then the thesis of unrestricted composition, moderate accounts of composition, and nihilism about composition are all on a par with respect to ontological parsimony. We can therefore base our theory choice on other criteria.

But what if composition is not identity? Is there some way of showing that mereology is ontologically innocent, other than by identifying wholes with parts?

In section 2 I will examine the idea that a commitment to the parts somehow automatically involves a commitment to the whole. On this view, accepting unrestricted composition does not bring any additional burdens because the burdens

were already incurred before the thesis of unrestricted composition was accepted. I call this the *levelling-up* account of ontological innocence: it entails that moderate and nihilist views of composition are just as costly as unrestricted composition, even if their advocates do not realise this.

In section 3 I will examine the idea that, although accepting unrestricted composition does expand your ontological commitments, this expansion is somehow not relevant to theory choice: these additions to the ontological burden do not affect the ‘price’ of the theory. I call this the *levelling-down* account of ontological innocence: it entails that unrestricted composition is no more costly, for purposes of theory choice, than are either moderate or nihilist views of composition. (The consequences for theory choice may seem the same either way, and we might wonder whether levelling up and levelling down are genuinely different options; I will return to this issue in section 4.)

The main goal of this paper is to investigate whether mereology can reasonably be called ‘ontologically innocent’ even if composition is not identity. This is an appealing prospect, for several reasons. First, it fits with much, though not all, of what Lewis says in section 3.6 of *Parts of Classes*, most obviously the reiteration of ontological innocence paired with the rejection of Donald Baxter’s strong composition as identity thesis. I do not think we can make full sense of everything Lewis says: some of his claims seem plausible only if composition is identity in a non-analogical way, but he explicitly denies the non-analogical thesis. (Sider 2007 and Bohn 2011 attempt to wrangle this.) Moreover Lewis clearly goes beyond claims of ontological innocence (pages 85-6 especially). And I do not think that there is a uniquely best way of vindicating the claims about ontological innocence without composition as identity. Indeed, some of Lewis’s remarks point towards the levelling-up account, and others towards the levelling-down account, whilst his remarks about Armstrong-style partial identity point towards a different notion of almost-innocence (compare Lewis 1993). Nevertheless, there’s much to be gained by examining the options.

Second, and relatedly, I think this captures what seems attractive in the passage I quoted at the outset: acknowledging the fusion of the cats doesn’t really add much to an ontology which already includes the cats; this isn’t a great extra ontological burden

to carry; it's not like adding the souls of the cats, the hive-mind of the cats, the invisible friends of the cats, or an extra eight legs for each cat. Can we substantiate these attractive thoughts without accepting composition as identity?

Third, if any argument in favour of unrestricted composition can be found in Lewis's discussion, it is an argument from the ontological innocence of composition claims, rather than an argument direct from the thesis that composition is identity. Ross Cameron has recently argued that composition as identity does not entail unrestricted composition (Cameron 2012; McDaniel 2010 has a different argument to the same conclusion). In outline, his argument is this: if composition is identity, then whenever a plurality has a sum, that sum just is its parts. But that fact doesn't determine the conditions under which a plurality has a sum, and in particular it doesn't entail that every plurality has a sum. 'Allowing that identity can be many-one simply doesn't tell us how *ubiquitous* cases of many-one identity are' (2012, 534, Cameron's italics). In other words, the General Composition Question is not the Special Composition Question (van Inwagen 1990, Hawley 2006, although compare Spencer 2012).

Nevertheless, composition as identity could provide defensive material for the thesis that composition is unrestricted, by entailing the ontological innocence of that thesis, thus rebutting the charge of ontological profligacy. If we can establish the ontological innocence of the thesis of unrestricted composition without resort to composition as identity, then this defensive move is still available. And this seems to be Lewis's approach: he does not infer unrestricted composition direct from composition as identity, but instead seems somewhat tempted towards composition as identity as a foundation for ontological innocence. So ontological innocence without composition as identity promises as much support for unrestricted composition as we should have expected from composition as identity itself.

2. Levelling Up and Ontological Commitment

In this section I investigate the levelling-up account of the ontological innocence of unrestricted composition. The idea is that a commitment to the existence of some objects automatically carries with it a commitment to the existence of their sum. Once you accept the existence of the cats, you are thereby committed to the existence of their fusion, whether you realise it or not. So accepting the thesis of unrestricted

composition does not give you any further commitments; the commitment to the existence of the fusion was already incurred when, as a small child, you incurred a commitment to the existence of kitties.

For the levelling-up account to have any plausibility, it must be possible for us to have ontological commitments we are unaware of: we can safely assume that Peter van Inwagen has examined his own beliefs, and yet he is adamant that he is committed to the existence of cats but not to the existence of cat-fusions (1990, 1994). Howard Peacock (2010) distinguishes the explicit from the implicit ontological commitments of a theory. Roughly speaking, the explicit commitments are what the theory says there is – the commitments which would be recognised by anyone who understood the theory – whilst the implicit commitments are those things which are required for the truth of the theory. If a commitment to the parts brings a commitment to the whole, as the levelling-up account has it, this must be an implicit rather than an explicit commitment. After all, van Inwagen understands his own beliefs yet does not recognise any commitment to cat fusions.

But what is it for something to be required by the truth of a theory? Peacock construes this modally: the requirements of the theory are those things which exist in every possible world in which the theory is true. Even setting aside issues about generic and specific requirements (Parsons 1970), this construal means that every possibly-true sentence carries ontological commitment to all necessary existents, such as numbers and perhaps God.

Whatever the independent merits of this modal construal of implicit ontological commitment, it cannot be what Lewis had in mind, because it does not discriminate between mereology and set theory. We are trying to underwrite the idea that a commitment to the parts just is a commitment to the whole, so that if you are already committed to the existence of the parts, then the claim that the whole exists is ontologically innocent. In contrast, a commitment to Possum is not supposed to involve a commitment to Possum's singleton: the claim that the singleton exists is ontologically non-innocent, according to Lewis. But the singleton of Possum exists in every world in which it is true that Possum exists, satisfying Peacock's modal account of implicit commitment.

We need a more fine-grained notion of implicit ontological commitment. Agustin Rayo (2007) takes the ontological commitments of a sentence to be the ontological demands which it imposes on the world; a sentence's ontological commitments are thus an aspect of its truth conditions, i.e. the demands it imposes on the world. As with Peacock's requirements, Rayo's demands reach beyond the explicit commitments immediately grasped by anyone who understands the sentence:

I shall use demand-talk in such a way that it obeys Kripke-style substitution rules for names and predicates. Thus, since Hesperus is Phosphorus, there is no difference between the demand that the world contain Hesperus and the demand that the world contain Phosphorus. Similarly, since being composed of water just is being composed of H₂O, there is no difference between the demand that human bodies be composed mostly of water and the demand that human bodies be composed mostly of H₂O. (Rayo 2007, 429)

However Rayo does not give a simple modal construal of demand-talk: not every object which must exist if the sentence is to be true counts as an ontological commitment of the sentence. The number two exists in every world in which 'St Andrews is in Fife' is true, yet it is not an ontological commitment of that sentence.

What then are the demands imposed on the world by a sentence's truth? Wisely, Rayo defers this question: 'The demands imposed on the world by a sentence's truth are simply the sentence's truth-conditions. So one's understanding of the former should be informed by one's understanding of the latter.' (2007, 429; see also Rayo 2008 section 3.2) Likewise, I will not attempt to explicate the notion of truth conditions here (though I touch on this again in section 2.3).

Given the levelling-up account, the thesis of unrestricted composition is ontologically innocent because the existence of the whole is demanded by the truth of the claim that the things which are its parts exist. For example, on this view it is an aspect of the truth conditions of 'Possum exists and Macavity exists' that the fusion of Possum and Macavity exists. In other words, ontological commitment transmits through the composition relation, just as it transmits through the identity relation. I will call this the *transmission thesis*. In the rest of this section, I explore the nature, justification, and dialectical role of the transmission thesis.

2.1 What is the Transmission Thesis?

There are weak and strong versions of the transmission thesis. On the weak version, if a sentence carries ontological commitment to some things, and they have a sum, then the sentence also carries ontological commitment to the sum. On the strong version, if a sentence carries ontological commitment to some things, then it also carries ontological commitment to their sum. If composition is in fact unrestricted, then the weak and strong transmission theses are extensionally equivalent: for every sentence, they will agree as to the ontological commitments of that sentence.

But the weak and strong transmission theses give different results if either moderate or nihilist views of composition are correct. For example, suppose that nihilism is true: there are no composite objects. Then, according to the weak transmission thesis, a commitment to the existence of some objects carries no further ontological commitment. But according to the strong transmission thesis, a commitment to the existence of some objects also carries commitment to the existence of their nonexistent sum. It is standardly thought that sentences can carry ontological commitment to nonexistent things; one very good reason for judging a sentence false is that it is ontologically committed to something which does not exist. This is why ontological commitment cannot be understood as a relation between sentences and things in the world.

Nevertheless, the strong transmission thesis is wildly implausible, even incoherent, if composition is not unrestricted. The ontological commitments of a sentence are the ontological demands imposed by the truth of the sentence upon the world. Suppose that Possum and Macavity each exist, yet they do not have a sum. The sentence 'Possum exists and Macavity exists' is true, but by hypothesis the sum does not exist. So evidently the truth of the sentence does not demand the existence of the sum. So the sum is not amongst the ontological commitments of the sentence, which is nevertheless committed to the existence of Possum and of Macavity. Therefore a commitment to the existence of Possum and of Macavity does not involve a commitment to the existence of their sum if they do not have a sum.

The levelling-up account requires a transmission thesis. The strong transmission thesis is certainly false if either a moderate or a nihilist view of composition is true. Moreover it is extensionally equivalent to the weak transmission thesis if composition is in fact unrestricted. To give the levelling-up account its best chance, I will therefore work with the weak transmission thesis.

Suppose the weak thesis is true. And suppose that van Inwagen (1990) is right that things have a sum when (and only when) they participate in a life together. If we are committed to the existence of some things, and if they participate in a life together, then we are committed to the existence of their sum. We may not know that we have a commitment to the sum (our commitment is not explicit), either because we are benighted nihilists who have not realised the truth of van Inwagen's moderate view, or because we are unaware of the empirical fact that those things participate in a life together.

Now continue to suppose that the weak transmission thesis is true, but suppose that van Inwagen is wrong and Lewis is right about composition. Then whenever we are committed to the existence of some things, we are committed to the existence of their sum. Again, we may fail to realise that we have this commitment, not for empirical reasons this time, but rather because we have not have realised the truth of Lewis's thesis of unrestricted composition.

2.2 Is the Weak Transmission Thesis True?

If composition is identity, the weak transmission thesis is true: implicit ontological commitment transmits through the identity relation, as between Hesperus and Phosphorus. (Composition as identity does not vindicate the strong transmission thesis. Ontological commitment transmits through the composition relation, if composition is identity, but this does not determine when the composition relation in fact holds. This illustrates, once again, the difference between the General and Special Composition Questions.)

Can we justify the weak transmission thesis without resorting to composition as identity? We can construct a notion of 'ontological fauxmitment' which transmits through the composition relation: if you have an ontological commitment to some

objects, and they have a sum, then by definition you have an ontological fauxmitment to the sum. Then our question is whether an ontological fauxmitment is a genuine ontological commitment.

There is no point in consulting criteria of ontological commitment, like Quine's famous dictum that a first-order sentence carries commitment to Fs just in case Fs must be counted amongst the values of the variables in order for the sentence to be true (for this formulation, see Rayo 2007, 430). Such criteria tell us about the circumstances under which sentences of certain kinds carry certain ontological commitments; they do not elucidate the notion of ontological commitment itself. Nor can we make progress by consulting our intuitive notion of ontological commitment, for there is no such intuitive notion (Sider 2011, 202).

Instead we need to pursue a larger project, considering the theoretical roles played by ontological commitment in the hope of establishing whether ontological fauxmitment also plays these roles. First, there is the connection with truth-conditions: the ontological commitments of a sentence are an aspect of its truth-conditions. Given that some objects have a sum, is it an aspect of the truth-conditions of the sentence which posits the existence of those objects that the sum exist also? Is the sentence in some sense *about* the whole, even though it does not explicitly mention the whole? (Compare: is 'Possum exists' in some sense *about* Possum's singleton, even though the set is not explicitly mentioned? Likewise: is 'the fans of the Chieftains are many' in some sense *about* the set of Chieftains fans?) Intuition does not guide me in this territory; a proper investigation would take us through debates about truth conditions and truth-making.

Second, there is a connection between ontological commitment and theory-choice. When we compare the ontological costs of different theories, should we include their ontological fauxmitments as well as their standard ontological commitments? Do the ontological fauxmitments contribute to the ontological profligacy of a theory, just as the standard ontological commitments do? If so, then it would be sensible to count ontological fauxmitments as genuine ontological commitments. I will discuss ontological profligacy and parsimony in section 3.

2.3 Dialectical Role of the Weak Transmission Thesis

If – *if* – the weak transmission thesis is true, if fauxmitments are genuine commitments, how should this affect the debate about theories of composition? If composition is in fact unrestricted, then accepting the truth of the unrestricted composition thesis brings us no additional ontological commitments: van Inwagen and the rest of us are already committed to the existence of arbitrary sums via our commitment to ordinary entities. If cat-fusions exist, then there is a sense in which we already believe in them, even whilst we strenuously deny that they exist. But if composition is in fact not unrestricted, then we do not already have these ontological commitments, so would be expanding the list of things to which we are ontologically committed if we mistakenly accepted that composition was unrestricted.

This creates a rather delicate situation regarding theory choice. Suppose that you accept the weak transmission thesis and are currently agnostic about the existence of any composite objects. You realise that if you follow Lewis and accept unrestricted composition you will believe in retrospect that your move was cost-free, and that your moderate and nihilist friends have failed to recognise the extent of their own commitments. You also realise that if instead you accept a moderate view of composition you will believe in retrospect that your move was cost-free, that your nihilist friend has failed to recognise the extent of her own commitments, and that your Lewisian friend's decision to accept unrestricted composition has committed her to the existence of many non-entities. Finally, you realise that if you opt for nihilism, you will think that both your moderate and your Lewisian friends have committed themselves to the existence of many non-entities.

You know what you will think in retrospect, whatever you decide. Whichever theory you accept, once you accept it you will believe that it was ontologically cost-free (this makes sense of Lewis's evangelical tone). But how should this affect your decision whether to accept that composition is unrestricted, if you are currently agnostic?

Set aside the other costs and benefits of the various theories of composition (for example, their interaction with debates about metaphysical indeterminacy); set aside questions about the relative badness of over-expanding versus over-contracting your ontology. If you want to use ontological parsimony as a criterion for theory-choice,

how should you apply it in this situation? There are two options. The first is to compare what the various theories would add to your current ontological commitments. The second is to compare what the various theories say about what the correct total ontology is.

The second option makes better sense, for two reasons. First, given the weak transmission thesis, you do not know what your ontological commitments are before you know which theory of composition is correct, so you cannot make the relevant comparison. Second, if we value ontological parsimony, then presumably this is because we think the world as a whole is likely to be well-represented by a parsimonious theory (*ceteris paribus*), not because we especially loathe *new* ontological commitments. If we value ontological parsimony, we should willingly accept a theory which brings us new ontological commitments, so long as it repudiates a greater number of our old ontological commitments. What matters is what a theory says about the total size of the world.

The thesis of unrestricted composition says that the world contains an enormous number of objects, many more than the world contains according to either moderate or nihilist theories of composition. So if we value ontological parsimony, this is a point against unrestricted composition.

The weak transmission thesis does justify the claim that, if composition is in fact unrestricted, then accepting the truth of the unrestricted composition thesis is an ontologically innocent move: we are already committed to the myriads of composite objects. So if we have good independent reason to think that composition is unrestricted, then we have good reason to claim that mereology is ontologically innocent, i.e. that accepting mereology merely makes explicit the vast ontological commitments we had implicitly incurred already. But this conditional fact is dialectically ineffective within the debate about composition; in particular it does not help the advocate of unrestricted composition to rebut the claim that her thesis is ontologically profligate.

3. Ontological Parsimony

According to the levelling-up account – discussed throughout section 2 – a commitment to some objects automatically carries a commitment to anything those objects compose, because ontological commitment transmits through the composition relation. This account levels up because it entails that, just as the advocates of unrestricted composition are committed to arbitrary sums, so too (unwittingly) are the advocates of moderate and nihilist accounts of composition, if composition is indeed unrestricted.

According to the levelling-down account – to be discussed in this section – accepting that composition is unrestricted involves commitment to extra entities, but this is irrelevant to theory choice. When we weigh up the costs and benefits of rival theories, these additional entities are cost-free, because they are related by composition to objects which are also acknowledged by rival views. This account levels down because it entails that, for the purposes of theory-choice, advocates of unrestricted composition are no worse-off in this respect than advocates of either moderate or nihilist accounts of composition, despite having more objects in their ontology.

The levelling-down account might be justified by talk of fundamentality or grounding. Suppose that wholes are grounded in or are less fundamental than their parts, or that only mereological atoms are fundamental. Suppose also that, for the purposes of assessing the ontological parsimony or profligacy of a theory, only the fundamental or ungrounded entities count. Then the levelling-down account of ontological innocence would be vindicated: unrestricted composition brings along extra entities, but not ones which really *matter*.

I will not pursue this line of justification in this paper, for several reasons. First, there are excellent discussions of these issues elsewhere (e.g. Cameron this volume, deRosset 2010, Schaffer 2008, Williams 2010). Second, this talk of fundamentality and grounding is insufficiently Lewisian for my present purposes. Third, I am trying to make sense of the alleged ontological innocence of mereology as following not from the metaphysical nature of the composition relation – whether that be a relation of identity or of grounding – but from primarily methodological or epistemic considerations.

3.1 Varieties of Parsimony

E.C. Barnes (2000) distinguishes *anti-quantity* from *anti-superfluity* principles of parsimony. An anti-quantity principle urges us to minimise the number of elements (individuals or kinds) in our theories. An anti-superfluity principle urges us to avoid elements which are superfluous by the lights of the theories they feature in. Suppose we are comparing two theories. One posits many entities, each with an explanatory role to play, whilst the other posits only a few entities, including one which is explanatorily idle. Prioritising anti-superfluity parsimony directs us towards the first theory, whilst prioritising anti-quantity parsimony directs us towards the second theory.

According to the levelling-down account, accepting the thesis of unrestricted composition commits us to the existence of many additional objects, but these objects do not ‘count’ for the purposes of assessing relative parsimony. The parsimony at stake here is best understood in terms of anti-quantity rather than anti-superfluity principles. This is because the additional objects – cat-fusions and worse – are not superfluous by the lights of the theory they feature in. The thesis of unrestricted composition does various jobs for Lewis and his fellow travellers, for example in solving the problem of the many, addressing puzzles about persistence, and avoiding both brutality and indeterminacy about composition. Most pertinently, the thesis is central to Lewis’s account of sets, which is why he dwells upon its alleged ontological innocence in *Parts of Classes*.

If arbitrary sums offend against a principle of parsimony, it is an anti-quantity principle. To assess whether they are truly offensive we need to make a further distinction amongst anti-quantity principles. *Quantitative parsimony* is determined by the sheer number of individual entities to which a theory is committed; *qualitative parsimony* is determined by the number of different kinds of entity to which a theory is committed. (Combining this standard terminology with Barnes’s useful distinction is somewhat confusing: anti-quantity principles urge us to favour both quantitative and qualitative parsimony.) I will consider these types of parsimony in turn.

3.2 Is Unrestricted Composition Quantitatively Profligate?

The thesis of unrestricted composition certainly has a *prima facie* problem of quantitative profligacy. The world according to unrestricted composition is much more heavily populated than the world according to moderate or nihilist views of composition; the levelling-down account of ontological innocence must explain why this population explosion does not count against the thesis of unrestricted composition.

In *Counterfactuals*, Lewis writes ‘I subscribe to the general view that qualitative parsimony is good in a philosophical or empirical hypothesis; but I recognise no presumption whatever in favour of quantitative parsimony’ (1973, 87). In context, this is convenient for Lewis, given his quantitatively-profligate ontology of possible worlds of the same qualitative kind as our actual world. It would help him here too, allowing him to disregard the quantitative profligacy of unrestricted composition. But Daniel Nolan (1997) and Alan Baker (2003) have persuasively argued that quantitative parsimony does matter, alongside qualitative parsimony and other theoretical virtues; no doubt opponents of unrestricted composition would agree with them. Moreover Lewis is making a point specifically about composition in section 3.6 of *Parts of Classes*, rather than dismissing quantitative parsimony in general.

So I will take it that the levelling-down account must explain why the quantitative profligacy of unrestricted composition does not count against it, given the assumption that quantitative parsimony counts in favour of a theory in other contexts.

Why does quantitative parsimony usually count in favour of a theory? According to Baker ‘Quantitative parsimony tends to bring with it greater explanatory power’ (2003, 258). Baker’s central example concerns the neutrino, which was postulated in order to explain why mass-energy and spin – quantities which are supposed to be conserved – seem to go missing when an electron is emitted from an atom during beta decay. Why did physicists hypothesise that a single neutrino with spin $\frac{1}{2}$ is emitted in each decay event, carrying all the ‘missing’ mass-energy and spin, instead of postulating that ten mini-neutrinos are emitted together, each with spin $\frac{1}{20}$ and each carrying $\frac{1}{10}$ of the ‘missing’ mass-energy?

The hypotheses are equally ranked for qualitative parsimony, each postulating one new type of particle (the neutrino versus the mini-neutrino). Both hypotheses are well-designed to explain the phenomena of beta decay, including the ‘missing’ mass-energy and spin. And, according to mini-neutrino theory, each mini-neutrino is causally active: none is a superfluous idle wheel. But, argues Baker, the mini-neutrino hypothesis fails to explain why, across various experimental scenarios, we never see entities of spin $1/20$ existing separately from the pack; nor does it explain why we never see entities with spin $7/20$, and so on. To explain these facts, we need to complicate the mini-neutrino hypothesis by stipulating that these particles inevitably hang out together in groups of ten. This in turn must either be added to the list of facts which need explanation, or else be added to the list of fundamental facts about the world. Either way, the overall picture looks messier than the single neutrino hypothesis. Quantitative parsimony is valued because it typically reduces explanatory complexity. (It is a further question why explanatory complexity is itself a theoretical vice; Huemer 2009 compares different rationales for preferring parsimony in science and philosophy.)

How might these considerations apply to parts and wholes? The thesis of unrestricted composition lacks quantitative parsimony. As well as cat-fusions and the like, it is committed to the existence of trout-turkeys, i.e. ‘the mereological fusion of the front half of a trout plus the back half of a turkey’ (Lewis 1991, 7). These extra entities certainly seem to generate additional explanatory complexity. Consider ‘congruence’ facts about the relationship between parts and wholes: we need to explain why trout-turkeys are always located in a sub-region of the region collectively occupied by a trout and a turkey, why they are inevitably somewhat scaly and somewhat feathery, why certain causal interactions with a trout constitute causal interactions with a trout-turkey, and so on. If there are no trout-turkeys, there are no such ‘facts’ to explain.

But we have explanations of such facts ready to hand, via our understanding of congruence facts for ordinary composite objects. If we can understand why turkeys are always located where their parts are, then we can understand why the same holds for trout-turkeys. And so on. So although unrestricted composition commits us to lots of extra entities, it does not require any additional explanatory complications in

order to explain how those entities behave, and how they relate to their parts. The usual reason for preferring quantitative parsimony does not hold sway in this case.

Nihilists about composition may still object (though see Sider 2003). Any theory which posits any composite objects, no matter how few, incurs explanatory burdens which the nihilist need not assume. (Nihilists have explanatory burdens of their own, of course.) So we may say that, for purposes of theory choice, the quantitative profligacy of the unrestricted composition thesis puts it at no disadvantage relative to moderate views of composition, although both unrestricted and moderate views are disadvantaged relative to nihilism.

3.3 Is Unrestricted Composition Qualitatively Profligate?

In very general terms, the additional composite objects countenanced by unrestricted composition are of the same kind as the entities accepted by moderate accounts of composition. Cat-fusions and trout-turkeys are just more composite material objects, with many familiar physical properties such as mass, charge, and shape. Moreover, they need not all have the peculiarly inflexible modal and temporal identity conditions often attributed to ‘mere sums’: advocates of unrestricted composition do not typically distinguish between ‘mere sums’ and other composite objects.

But in less general terms very many of these additional objects seem to fall under novel kinds: they add qualitative diversity to our ontology, not just quantitative diversity. The extent of this additional qualitative diversity will depend upon the rival theories to which unrestricted composition is compared. For example, van Inwagen, who believes that all composite objects are alive, would point out that unrestricted composition commits us to the existence of a significantly different new kind of thing: inanimate composite objects. Even those who accept the existence of armchairs, buildings and continents would point out that unrestricted composition brings commitment to some very peculiar new species such as the trout-turkey

We can distinguish three sorts of concern about trout-turkeys. The first is quantitative profligacy, as already discussed: the trout-turkeys are yet more objects. The second is qualitative profligacy: given that the kind *trout-turkey* is supposed to supplement, not replace, the kinds *trout* and *turkey*, this inflates the number of kinds in our ontology.

The third is a more direct objection to trout-turkeys (and other arbitrary sums) *per se* on the grounds that they are strange entities, not recognised by common sense.

Whatever the merits of this third objection, it is not addressed by the levelling-down account of ontological innocence. If you have a direct intuition that there simply is nothing composed of the front half of this trout and the back half of that turkey, then you will not be consoled by the thought that, although there is such a thing, it does not render unrestricted composition problematically profligate. You might be consoled by the levelling-up account, according to which you are already committed to trout-turkeys despite your intuitions, but I doubt it. You might be consoled by the claim that composition is identity and that trout-turkeys are therefore nothing to worry about, but (Cameron 2012: 551) will then show you the errors of your ways. The third trout-turkey concern – the objection from common sense – lies beyond the scope of this paper, which is concerned with questions of ontological innocence.

So the levelling-down account needs to explain why it is that, although the thesis of unrestricted composition commits us to the existence of new kinds of thing, this does not count against the thesis when we compare it to rival theories of composition. Why does qualitative parsimony usually count in favour of a theory? As with quantitative parsimony, explanatory complexity seems key. A theory which posits more kinds of entity must explain the relationships between these different kinds, as well as explaining the behaviour of individual instances of the various kinds.

As with quantitative profligacy, the levelling-down account has some plausibility here. The behaviour of trout-turkeys, cat-fusions and other arbitrary sums is correlated with the behaviour of more familiar objects in predictable ways. We need not invent new (biological?) laws to govern trout-turkeys, nor, of course, regard the property of being a trout-turkey as natural to any high degree. We might instead start to worry that these kinds are superfluous, since they seem explanatorily redundant. But the advocate of unrestricted composition can accept this concern, and deny that *trout-turkey* is a genuine kind. This does not render the individual trout-turkeys superfluous in the context of the theory; as I argued above, arbitrary sums do significant philosophical work for Lewis.

Again, the levelling-down account succeeds in reducing the cost (in this respect) of unrestricted composition so that it is equal to the cost of moderate theories of composition: once we are committed to the existence of some composite object, commitment to further composite objects does not generate additional explanatory complexity. But the account does not manage to reduce the cost to that of bargain-basement nihilism. Any non-nihilist account of composition is at a disadvantage with respect to ontological parsimony; non-nihilist accounts may still be preferable on other grounds, however.

4. So, *Is Mereology Ontologically Innocent?*

We have investigated whether mereology can reasonably be called ‘ontologically innocent’, if composition is not identity. If ontological commitment transmits through the composition relation – if fauxmitments are commitments – then we are all already committed to the existence of whichever composite objects are countenanced by the true theory of composition, whether that be nihilism, a moderate view, or the thesis that composition is unrestricted. But this doesn’t help make the case for unrestricted composition and its ontologically-profligate picture of the world.

The more promising strategy is to argue that, once we have accepted the existence of at least some composite objects, the addition of more and weirder such objects to our ontology does not create additional explanatory burdens, so does not count against unrestricted composition as compared to moderate views of composition.

When I introduced the levelling-up and levelling-down accounts in section 1, I acknowledged their apparent similarity. In hindsight, however, they seem quite different. According to the levelling-up account, we all have the same ontological commitments (assuming a shared stock of simples), though we do not have any theory-neutral way of working out what in fact we are committed to. We can, however, all see that the thesis of unrestricted composition paints a more ontologically-profligate picture of the world. According to the levelling-down account, different theorists of composition have different ontological commitments, but this does not count against the thesis of unrestricted composition as opposed to moderate views of composition.

Could we somehow combine the levelling-up and levelling-down accounts? Suppose fauxmitments are commitments, so we are all already committed to the existence of those composite objects countenanced by the true theory of composition, whatever that is. Could we nevertheless argue that unrestricted composition is at least as plausible as moderate views, because it is explanatorily no more complex than moderate views? Accepting the levelling-up account does not prevent us from also accepting the levelling-down account, but the combination of these accounts does not seem to lend any stronger support to the thesis of unrestricted composition than does levelling-down alone.

5. Is Decomposition Ontologically Innocent?

I have been discussing rival answers to the Special Composition Question, which asks about the conditions under which some things compose a whole. The Inverse Special Composition Question asks about the conditions under which a thing has proper parts; equivalently, the Simple Question asks about the conditions under which an object is mereologically simple (van Inwagen 1990, Markosian 1998). Call an answer to the Inverse SCQ a ‘decomposition thesis’. Are decomposition theses ontologically innocent? That is, once you are committed to the existence of an object, is it ontologically innocent to accept that the object has proper parts?

This suggestion might seem puzzling. If both composition and decomposition are ontologically innocent, how do we ever incur substantive ontological commitments? But recall that both composition theses and decomposition theses typically make only conditional existential claims: if there are some cats, they have a fusion; if there is an extended object, it has a part in each of ‘its’ proper sub-regions. Substantive commitment comes when we accept the antecedent of such a conditional, and our question is whether accepting the consequent is then an ontologically innocent move.

Some of Lewis’s remarks suggest that ontological innocence works both ways: ‘It just *is* them. They just *are* it...Commit yourself to their existence all together or one at a time, it’s the same commitment either way’. Moreover identity is symmetric, so if composition is identity we should expect consequences for decomposition theses if there are consequences for composition theses (deRosset 2010, Spencer 2012).

According to a levelling-up account of the ontological innocence of decomposition, if the truth of a sentence requires the existence of some object, then it requires the existence of any parts that object happens to have ('requires' here is non-modal). If you are committed to the existence of an object, then you are implicitly committed to the existence of whatever parts it has, whether you realise this or not. As with composition, such implicit commitments cannot play a significant role in debate about decomposition, nor can they address concerns about ontological profligacy.

According to a levelling-down account of the ontological innocence of decomposition, once we are committed to the whole, adding the parts does not increase the ontological 'cost' of the theory in question, because they do not increase its explanatory complexity. As with composition, we could base this claim on the metaphysical relationship between parts and wholes: perhaps parts are grounded in wholes, and only ungrounded entities 'count' for the purposes of measuring ontological parsimony (Schaffer 2007, 2010). Again, however, I am trying to establish ontological innocence without recourse to the metaphysics of the (de)composition relation. Moreover, basing ontological innocence on notions of grounding means that either composition or decomposition theses may be innocent, but not both.

How else might we argue for the levelling-down account? As with composition, once we have an explanatory framework based on the relation between parts and wholes, adding extra entities to the picture does make much difference. If we have a place for a particular object in our ontology, plus an understanding of why in general parts are located in sub-regions of the whole, then we do not create any further explanatory tasks by accepting that this particular object has proper parts. So if there are *some* proper parts, then the thesis that there are *more* proper parts is ontologically innocent. Nihilists about composition still have the upper hand in this respect: whether they countenance teeming masses of tiny simples, or one big simple universe (Schaffer 2007), adding parthood relations to this picture increases its complexity.

References

Bohn, Einar (2011): 'Commentary on David Lewis *Parts of Classes*', *Humana.Mente Journal of Philosophical Studies*, 19: 151-8.

- Cameron, Ross (2012): 'Composition as Identity Doesn't Solve the Special Composition Question', *Philosophy and Phenomenological Research* 84.3: 531-54.
- Cameron, Ross (this volume): 'Parts Generate the Whole, but They are Not Identical to It', *Composition as Identity*, edited by Donald Baxter and Aaron Cotnoir, Oxford: Oxford University Press.
- deRosset, Louis (2010): 'Getting Priority Straight', *Philosophical Studies* 149:73-97.
- Hawley, Katherine (2006): 'Principles of Composition and Criteria of Identity', *Australasian Journal of Philosophy* 84.4: 481-93.
- Huemer, Michael (2009): 'When is Parsimony a Virtue?' *Philosophical Quarterly* 59: 216-36.
- Lewis, David and Lewis, Stephanie (1970): 'Holes', *Australasian Journal of Philosophy*, 48: 206-12.
- Lewis, David (1991): *Parts of Classes*, Oxford: Blackwell.
- Lewis, David (1993): 'Many But Almost One', in *Ontology, Causality and Mind*, edited by J. Bacon, K. Campbell and L. Reinhardt, Cambridge: Cambridge University Press 23-38.
- Markosian, Ned (1998): 'Simples', *Australasian Journal of Philosophy* 76.2: 213-28.
- McDaniel, Kris (2010): 'Composition as Identity Does Not Entail Universalism', *Erkenntnis* 73: 97-100.
- Merricks, Trenton (2001): *Objects and Persons*, New York: Oxford University Press.
- Nolan, Daniel (2005): *David Lewis*, Chesham: Acumen.
- Nolan, Daniel (forthcoming): 'Lewis's Philosophical Method', in Loewer and Schaffer (eds.) *Blackwell Companion to David Lewis*, Oxford: Blackwell.
- Parsons, Terence (1970): 'Various Extensional Notions of Ontological Commitment', *Philosophical Studies* XXI: 65-74.
- Peacock, Howard (2011): 'Two Kinds of Ontological Commitment' *Philosophical Quarterly*, 61: 79-104.
- Rayo, Agustin (2007): 'Ontological Commitment', *Philosophy Compass* 2/3: 428-44.
- Rayo, Agustin (2008): 'On Specifying Truth-Conditions', *Philosophical Review*, 117: 385-443.
- Schaffer, Jonathan (2007): 'From Nihilism to Monism', *Australasian Journal of Philosophy* 85.2: 175-91.

- Schaffer, Jonathan (2008): 'Truthmaker Commitments', *Philosophical Studies* 141: 7-19.
- Schaffer, Jonathan (2010): 'Monism: The Priority of the Whole', *Philosophical Review* 119.1: 31-76.
- Sider, Theodore (2003): 'What's So Bad about Overdetermination?' *Philosophy and Phenomenological Research*, 67: 719-26.
- Sider, Theodore (2007): 'Parthood', *Philosophical Review*, 116: 51-91.
- Sider, Theodore (2011): *Writing the Book of the World*, Oxford: Oxford University Press.
- Spencer, Joshua (2012): 'Strong Composition as Identity and Simplicity', *Erkenntnis*, published online 20th July 2012.
- Van Inwagen, Peter (1990): *Material Beings*, Ithaca NY: Cornell University Press.
- Van Inwagen, Peter (1994): 'Composition as Identity', *Philosophical Perspectives* 8: 207-20.
- Williams, J. Robert G. (2010): 'Fundamental and Derivative Truths', *Mind* 119: 103-141.