



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/92249/>

Version: Accepted Version

Article:

Caves, RLJ (2018) Emergence for Nihilists. *Pacific Philosophical Quarterly*, 99 (1). pp. 2-28. ISSN: 0279-0750

<https://doi.org/10.1111/papq.12148>

© 2018, The Author. *Pacific Philosophical Quarterly*. © 2015, University of Southern California and John Wiley & Sons Ltd. This is the peer reviewed version of the following article: "Caves, R. L. J. (2015) Emergence for Nihilists. *Pacific Philosophical Quarterly*" which has been published in final form at <http://dx.doi.org/10.1111/papq.12148>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Self-Archiving.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Emergence for Nihilists

Richard L. J. Caves

Abstract: I defend mereological nihilism, the view that there are no composite objects, against a challenge from ontological emergence, the view that some things have properties that are ‘something over and above’ the properties of their parts. As the nihilist does not believe in composite wholes, there is nothing in the nihilist’s ontology to instantiate emergent properties – or so the challenge goes. However, I argue that some simples (taken together) can collectively instantiate an emergent property, so the nihilist’s ontology can in fact accommodate emergent properties. Furthermore, I show that employing plural instantiation does not bloat the nihilist’s ontology or ideology.

1. *The Argument from Ontological Emergence*

Here I defend mereological nihilism against a challenge from ontological emergence. §1 briefly explains nihilism and ontological emergence then outlines the argument against nihilism from the existence of emergent properties (the No Bearers argument). In §2 I argue that the best solution is for the nihilist to employ plural instantiation, holding that emergent properties are collectively instantiated by multiple simples. §3 weighs up the ontological and ideological cost of introducing plural quantification – necessary for making sense of plural instantiation – into the nihilist’s theory. I conclude that the cost is negligible and so does not negate any of the claimed ontological or ideological parsimony advantages of nihilism. §4 defends and elaborates my views on ideological parsimony employed in the previous section, while §5 clarifies how my Plural Instantiation Strategy advances the debate and defends the strategy against an objection from the unity of consciousness.

Suppose we have good reason to believe that wholes have novel properties over and above the properties of their parts. Next suppose – in a drive for ontological austerity, perhaps – that we get rid of the wholes from our theory: can we find a way to coherently continue to believe in the novel properties we once attributed to them? I’m going to try and convince you that the prospects for doing so are not as dim as they might seem at first glance. That is, I am going to defend Mereological Nihilism (nihilism) – my favoured answer to the Special Composition Question (SCQ) – against an argument from the possibility of Ontological Emergence (emergence).

SCQ: Under what circumstances, when we have some things, do they compose some further thing?¹

There are numerous potential answers to SCQ. Most answers place some non-trivial restriction on when composition occurs and when it does not. These *restrictivist* answers range from naïve answers such as ‘when the objects are touching’ or ‘when the objects are fastened

together’, to more considered answers such as ‘when the objects participate in a life’.² There are also two *extreme* answers to SCQ: Mereological Universalism places no restriction on composition. The universalist answers that whenever there are some things they *always* compose some further thing.³ The nihilist, on the other hand, says that when you have some things they *never* compose some further thing. As such, nihilism entails there are no composite objects: that when you have some things (maybe some subatomic particles), they never compose some further thing (such as a cat). Rather, there are just mereological simples, arranged in a variety of ways (cat-wise, for instance).⁴ Potential reasons for favouring nihilism include its ability to provide a uniform eliminative answer to a number of puzzles involving composite objects⁵ and its ability to provide us with a parsimonious ontology and ideology.⁶ However, my purpose here is not to argue that nihilism is the best answer to SCQ, but merely to defend it from the argument from emergence. Nihilists and believers in composition differ over whether or not composite objects exist, yet we have said nothing concerning the nature of composite objects— what such objects are like, if they exist. One can, in principle, defend an answer to SCQ while remaining silent on this further question: but, as we’ll see, the argument on emergence puts pressure on interlocutors on both sides of the mereology debate to say something informative here.

In particular, we can ask: are the properties of a whole necessarily derived from the intrinsic properties of its parts (and the spatial relations between them)? For example, that the egg is partly white and partially yellow (and the yellow part is inside the white part) would seem to be straightforwardly derived from it having a yellow part (the yoke) and a white part (and the yoke being located inside the white). Or might some of the whole’s properties not be derivative in this way— might they be something *over and above* these derivative properties? According to the believer in ontological emergence, it is at least sometimes the case that when you have some things and you arrange them in a certain way, you not only get some further object, but some further property – an *emergent property* – instantiated by that further thing. Emergent properties are supposed to be something genuinely in addition to; not metaphysically reducible to; the properties of the things you started out with: the whole is greater than the sum of its parts, having at least one property that is neither inherited from those parts, nor somehow derivative of the properties of those parts.⁷ More carefully: we should understand emergent properties as fundamental (non-derivative) properties which are distinct from, but are none the less dependent on, a set of fundamental microphysical base properties.⁸ Individual base properties are understood as instantiated by individual mereological simples, while any wholes composed of these simples will have properties that depend on these base properties: the question of ontological emergence is then whether any of these dependent properties are fundamental.⁹

The most often discussed (putative) example of this phenomenon hails from the Philosophy of Mind: when you have some neurons and arrange them in a certain very special way – brain-wise, that is – not only do you have some further object, a brain, the thought goes, but some further property (or properties) – namely, phenomenal consciousness (or qualia) – that is something genuinely *over and above*, or not metaphysically reducible to, all the properties of the individual neurons and the spatial relations between them taken together. Another important example is from Quantum Mechanics.¹⁰ Schaffer (2010), for instance, argues that an entangled quantum system has properties, such as spin, not reducible to the properties of its parts and the spatial relations between them: For an entangled electron pair, Schaffer says, ‘a pure spin state can be attributed to neither electron individually. A pure spin state can be attributed to the electron pair only collectively, as a system’ (p.51). Given that the spatial relations between the particles are irrelevant to their collective spin, and treating the ‘system’ as the mereological sum of the two particles, we thus have an example a whole with properties genuinely over and above the properties of its parts.¹¹

The *prima facie* challenge to the nihilist is fairly intuitive: if you don't believe in composite objects, how can you account for the existence of emergent properties? Assuming we have independent reason to believe in the existence of emergent properties, we have reason to posit composite objects as their bearers. For suppose a believer in composite objects grants, for sake of argument, that we have no evidence for the existence of chairs over and above the evidence we have for believing in simples arranged chair-wise: That there appears to be a four-legged wooden object in region *R*, and that when I sit down in region *R* I don't fall on the floor, is just evidence for there being simples arranged chair-wise in region *R* – it doesn't tell between there being some further object, a chair, in region *R* over and above the simples so arranged, on the one hand, and their being *merely* some simples arranged chair-wise on the other. However, our believer might go on to say, the fact that I have (say) the subjective experience of diverse yet unified phenomenal feels – i.e. the existence of emergent consciousness – is evidence for the existence of some thing (a mind or brain) over and above the simples arranged brain-wise in the vicinity of my head.

The crucial thought here is that emergent properties, such as consciousness, require the existence of a bearer, and the nihilist doesn't provide any viable candidates to be *suitable* bearers of such properties. The nihilist only believes in simples: tiny (perhaps point-sized) mereological atoms that – very plausibly – lack the rich internal structure needed to be, for instance, thinking, feeling beings. One place where this issue plays out prominently in the literature is in relation to the respective cases presented by van Inwagen (1990) and Merricks (2001) for the systematic elimination of composite objects: both would entail nihilism, if not for the fact that both arguments allow that if a composite object has a property which is not reducible to the properties of its parts, i.e. an emergent property, then it is exempt from elimination. Van Inwagen (1990, p.118) argues that thinking is an activity that must be done by a single composite whole (such as a human being), since thinking cannot be done by many smaller things in the way many atoms can be arranged to perform the function of a shelf without actually composing a shelf (and no individual simple can think).¹² Merricks (2001), meanwhile, is explicit that consciousness has powers not derived from its microphysical base, rejecting pure nihilism because of a belief that a given human being, by virtue of being conscious, possesses at least one property 'not, of metaphysical necessity, implied by the existence and intrinsic properties of, and spatiotemporal and causal interrelations among [its] constituent atoms.' (p.89) Here we have a move from emergent consciousness to the failure of nihilism.

Now, both van Inwagen and Merricks give reasons for believing that we ought to avoid eliminating *us* from our theories if at all possible, which I won't go into here. However, suppose you were impressed by van Inwagen or Merrick's general case for the elimination of composite objects and weren't actively looking for a reason to keep any such objects around: it would seem natural to try to move towards a more unified theory by getting rid of us along with everything else. What would stand in your way? Well, if we grant that consciousness is a property that can't merely be explained by the intrinsic properties of individual simples and the external relations between them, then it seems what we're lacking is a suitable *bearer* for consciousness to, as it were, get the thinking done. *Ex hypothesi*, neither some simples being merely arranged thus-and-so, nor any intrinsic property of any individual simple, can be numerically identical with any of the instances of consciousness we take to be instantiated in the world: but at first sight this exhausts all the options for the nihilist. Similarly, Schaffer (2007) argues that the nihilist cannot accommodate emergent properties – at least not if she believes in a plurality of simples – since wholes have spin properties that are not derivative from the spin of their parts and the (spatial, causal) external relations between them.¹³ So, we have fundamental holistic spin properties that need to be accommodated and no individual simple is a suitable bearer for the collective spin of more than one simple: the conclusion reached is that, with only simples in her ontology, the nihilist must say there is only one simple (the World) so that the only fundamental spin property

will be the intrinsic spin of that one simple. So, is pluralistic nihilism really incapable of accommodating these kinds of properties?

Before moving on to potential nihilist solutions, let's give a more precise formulation of the argument against nihilism from emergence:

The 'No Bearers' argument –

- P1. If a property is instantiated in the world, there exists a suitable bearer.
- P2. There are emergent properties instantiated in the world.
- P3. There exist suitable bearers for emergent properties (from P1 + P2).
- P4. If nihilism is true, the only possible bearers of any properties are mereological simples.
- P5. Mereological simples are not suitable bearers for emergent properties.
- P6. If nihilism is true, there are no suitable bearers for emergent properties (from P4 + P5).
- C1. *Therefore:* Nihilism is false (from P3 + P6).

Potential replies to No Bearers on behalf of the nihilist will fall broadly into two camps: On the first, grant for sake of argument the assumption that there are emergent properties and show that the nihilist can provide suitable bearers for them (rejecting P6 via P5). In defending such lines of response, the nihilist need not be concerned, dialectally, with trying to provide their own detailed account of emergent properties or giving reason to think they are the best way explaining the target phenomenon: the nihilist will be able to say whatever her opponent says about emergent properties and why they're applicable to the target phenomena. If her opponent is doomed to fail in this task, the nihilist is more or less off the hook: If P2 is false, there is no direct argument against nihilism from the existence of putatively emergent phenomena. The other style of response will deny that emergent properties provide the best account of putatively emergent phenomena such as phenomenal consciousness or quantum entanglement. Once we have an alternative account, it will be argued, we'll see that we already have a ready explanation of the problematic phenomena, and so positing emergent properties would be gratuitous. If there are no emergent properties, there need not exist suitable bearers (undercutting support for P3 via the denial of P2).

Examples of this latter sort of response in the quantum mechanical case include the strategy of positing entanglement relations, and employing the ideology of metaphysical indeterminacy as an ontologically innocent stand-in for entanglement relations.¹⁴ In the Philosophy of Mind, any theory which seeks to explain consciousness as anything other than an emergent property – e.g. eliminativism, reductive physicalism, Cartesian dualism – falls into this camp. No such strategy will be discussed in detail in this paper. Instead, we will be interested in the question of whether, if we grant for sake of argument that there are emergent properties, the nihilist can accommodate them in her ontology. That is, whether suitable non-composite bearers can be found.

In the search for a non-composite bearer for emergent properties, this paper grants for the sake of argument the assumption that a lone simple is not a suitable bearer for an emergent property. It's worth flagging up, however, that there are some interesting options worth considering before giving up on lone simples as bearers of emergent properties. After all, individual *microphysical particles* aren't suitable bearers for properties such as phenomenal

consciousness or entangled spin properties, but why assume all simples are microphysical particles? For instance, why couldn't whatever object in fact instantiates my consciousness – which I take for granted isn't a microphysical particle – be mereologically simple? The most persuasive objection, I think, is that simples lack the requisite internal structure to instantiate a rich, complicated macro-property such as consciousness, which is presumed to be the product of the interactions of countless micro-particles in very particular arrangements.¹⁵ That is, you might think that simples necessarily lack internal heterogeneity by virtue of their simplicity: a simple can't have a positive electric charge *here*, and a negative electrical charge *there*, for instance, because this requires having some proper part that is positively charged *simpliciter*, and some distinct proper part that is negatively charged *simpliciter*, and simples by definition have no proper parts (see McDaniel 2007). However, there have been some compelling arguments put forward in the literature to show that simples could be heterogeneous and spatially extended by possessing fundamental distributional properties: e.g. the property of being positively-charged-here-and-negatively-charged-there, where this property is not had by virtue of it or its parts possessing any non-distributional charge properties.¹⁶ If this is right – if mereologically simple things could have limitless internal heterogeneity by instantiating irreducible distributional properties – the main objection to simples not being suitable bearers for emergent properties is undercut.

So, could the object (my brain?) that instantiates my consciousness be a heterogeneous extended simple? One that emerges from the microphysical base in much the same way as my consciousness is being supposed to emerge, as effectively a purpose-built bearer for my emergent consciousness? Although this may seem outlandish, I would need convincing that – when the proposal is properly chewed over and digested – it is any more strange or counterintuitive than the emergence of the property of consciousness itself. However, although I think this option warrants a full discussion, I'm cutting it short to explore another possibility I take to be promising: plural instantiation. This strategy will be explained in the next section and be our focus for the remainder of the paper.¹⁷

2. The Plural Instantiation Strategy

Perhaps numerous simples, acting together, could instantiate an emergent property? That is, perhaps there are emergent property instances that, instead of having a single bearer, have multiple bearers, which jointly bring about the conditions necessary for the emergent property to be instantiated. Let's call this plural (or collective) instantiation. As I'll be arguing, plural instantiation is a possibility we should take seriously. Given plural instantiation, I see no reason why a plurality of simples could not jointly give rise to a single instance of a given emergent property. The factors that might be taken to disqualify individual simples from being suitable bearers of emergent properties are not present when we consider a plurality of such entities: The individual simple is not (let's grant) spatially extended, but the plurality is dispersed across the same region of space that any emergent property they collectively instantiate would be extended over. An unextended simple cannot be qualitatively heterogeneous, but the plurality can be simply in virtue of different members of the plurality possessing differing (qualitatively homogeneous) properties.

The nihilist's opponent assumes that that when we have a case of emergence, we have some simples that give rise to *some further particular* that instantiates the emergent property in question.¹⁸ On the Plural Instantiation strategy, by contrast, no further entity is posited – there are just the simples, which jointly instantiate the emergent property. They cut out the middleman, so to speak. Instantiation here is not taken to be a one-one relation between some object 'the plurality'

and the emergent property being instantiated; rather it is a many-one relation between the many simples, on the one hand, and the emergent property on the other: we're not employing run of the mill one-one instantiation here. This is what allows the proponent of this strategy to posit some simples arranged brain-wise, say, that together instantiate the emergent property 'being conscious', without thereby composing some further object, a brain. It's also what makes this strategy somewhat revisionary. Indeed, this whole approach hinges on the plausibility of introducing plural instantiation. Do I have an argument for plural instantiation? Only that there is a dearth of arguments against it and no obvious cost involved in employing it in our theorizing. In what follows we'll try to get clear on what exactly the view entails and I'll try to deal with any potential worries or misunderstandings that arise along the way.

From the outset plural instantiation should be distinguished from multiple instantiation.¹⁹ A property is multiply instantiated if it has many separate instances: Red is multiply instantiated since there are many red things. Each of these red things is a separate instance of redness. Plural instantiation, on the other hand, occurs when there is more than one object contributing to the very same instance of a property. For example, the fleas collectively engulf the cat, but no flea individually engulfs him: there is just one instance of the property of engulfing a cat in the vicinity of the cat. Call any property which is plurally instantiated a plural property. A property that is not plural is singular. It might help here to distinguish a collective predicate from a distributive predicate: if a distributive predicate F applies to some things then those things are F because each of those things is individually F. However, if a collective predicate G applies to some things then those things are G but none of those things is, by itself, G. *If our language is wearing its metaphysics on its sleeve*, then distributive predicates will predicate singular properties of things, while collective predicates will predicate plural properties of things. Being a plural property is not the same thing as being a relation, even though both involve many things participating in bringing about the one property instance. As we shall see, the *monadic property / relation* distinction crosscuts the *plural / singular* distinction, allowing for monadic plural properties and singular relations. For example, if the chairs are three metres away from the stage, it is true of each of the chairs individually that they are three metres away from the stage: this is a straightforward singular relation (and 'three metres from' is a distributive, as opposed to a collective, predicate).

On the other hand, consider the following couple of examples: First, the citizens of a democracy collectively enjoy sovereignty, while no individual citizen is sovereign. Intuitively, although we might sometimes speak of the 'The People' as if it were some singular composite entity that alone enjoys sovereignty, we want to say that it's the individual people considered jointly that have sovereignty, not some hypothetical mereological sum of those people. Second, Jess looks up at the night sky and declares that the stars are beautiful: Asked if any individual star is beautiful on its own, she's inclined to respond in the negative – sure, an individual star can be quite pretty, but it is after all just a tiny pin prick of light. Asked if it's the *arrangement* of the stars which is beautiful, she's also inclined to answer negatively: their being arranged in a certain way in the sky contributes to *the stars themselves* being beautiful, but the arrangement itself would be nothing special if the stars were substituted for salt crystal or tea leaves. No, according to Jess it's the stars that – jointly – instantiate the monadic property of beauty. It doesn't seem like Jess's opinion is crazy or incoherent.

You might quibble about whether sovereignty or beauty are genuinely monadic properties. What we have in each case is a monadic, collective predicate— but plausibly in these cases our language *doesn't* wear its metaphysics on its sleeve, and there is no genuinely plural property being predicated. Certainly, unless we think the properties in question are emergent, these are not *fundamental* properties corresponding to the collective predicate: thus we will say that such properties are ultimately grounded in the parts and the external relations between them after all.

But whatever the ultimate metaphysical structure of the cases, we can make sense of ‘is sovereign’ and ‘is beautiful’ as being monadic, collective predicates in ordinary language: the claim that there are plural properties is just the claim that at least some of the time these sorts of predications correspond directly with the underlying metaphysics. That there is any place for plural instantiation in serious metaphysics, or fundamental reality, is precisely what I’m arguing in the paper as a whole, with respect to emergent properties; but you should at least be convinced that the idea of a jointly instantiated monadic property is coherent.

To help get clearer on the various distinctions at work, we can consider the number of argument places had by a predicate standing for a given property or relation. Predicates of monadic properties have just one argument place: For example ‘___ is beautiful’ has one space to be filled. Whereas a two-placed relation, such as ‘___ loves ___’ has two spaces to be filled. A singular property or relation has each space filled by a single object: for instance, ‘The Morning Star is beautiful’ and ‘Troilus loves Cressida’, respectively. A plural property or relation, on the other hand, has at least one argument place filled by more than one object. For example, ‘the (individually unremarkable) stars are (collectively) beautiful’ and ‘the (individually tiny) fleas (collectively) engulf the cat’. The intuitive point is that each argument place in a relation assigns a distinctive role to the object or objects that fill that place: while a case of plural instantiation is where many objects collaborate to play the one role and so collectively fill the one argument place. To reiterate, all the putative examples of plural instantiation above have been provided just to help the reader get a fix on what plural instantiation is supposed to be: it doesn’t matter for these purposes if you don’t agree that any of these are genuine cases of plural instantiation, provided you understand what is being claimed. It seems fairly obvious to me that the English language has irreducibly collective predicates such as ‘surround’ – and, indeed, ‘compose’ – but that’s also orthogonal to the debate here. What’s at issue is whether the nihilist is free to introduce plurally instantiated properties into her ontology in order to meet the challenge from emergence.

Given all this, I see no conceptual barrier to the nihilist holding that the many simples are the collective bearers of a single, monadic emergent property. You might worry that including plural properties in your ontology requires significant revision to your logic. However, as Byeong-Uk Yi (1999, 2002, 2005, 2006) and others²⁰ have demonstrated, we can accommodate plural properties simply with the introduction of plural quantifiers and plural predicates to elementary logic (PFO+). Given the relative ease with which such predicates can be accommodated, I don’t think there’s any sustainable objection to effect that the very idea plural properties is confused, or that formulating a theory in a language that contains plural predicates would be overly arduous or unattractively revisionary. That is, I don’t think a metaphysics containing fundamentally collective properties or plural instantiation can be dismissed out of hand.²¹

Having argued for the coherence of the plural instantiation strategy, I will us the next two sections to argue that implementing the strategy does not come at too high a theoretical cost, while in the final section I defend the strategy from a worry that consciousness cannot be plurally instantiated and so is out of the reach of my general strategy— this latter discussion also enables me to clarify how my defence of the Plural Instantiation represents a novel contribution to the debate. If everything I say in this section and the subsequent section goes through, I think we have enough to conclude that the Plural Instantiation strategy is a serious and attractive option for the nihilist who wishes to accommodate ontological emergence.

3. The Cost of Plural Instantiation

We've seen that the nihilist can accommodate emergent properties in her theory if they are plurally instantiated properties. But just because the nihilist is at liberty to introduce plural properties into her theory without falling into conceptual confusion or straying beyond the philosophical pale, doesn't necessarily mean it would be prudent to do so. Such a move may still come at a cost: the price of introducing plural instantiation into your theory may be one that the nihilist would be unwilling, or unwise, to pay. How might that price be exacted? Plausibly through reducing the ontological or ideological simplicity of the nihilist's theory. We'll assess each possibility in turn. In each case, we'll see that at the centre of the controversy is the strategy's reliance on plural quantification. The nihilist needs to be able to quantify plurally over *some things* in order to say of them that they satisfy a given collective predicate and thus plurally instantiate an emergent property. For, *ex hypothesis*, it's not the case that any individual thing instantiates an emergent property, and so it's not true of any individual thing that might be picked out by a singular existential quantifier that it satisfies a collective predicate.

Many have worried that plural quantification is not ontologically innocent.²² The worry is that our best metaphysical theory is not only ontologically committed to each thing in the domain of its singular existential quantifier, but also to every 'plurality' in the domain of its plural quantifier, should the theory contain one. The thought is we should extend the Quinean criterion of ontological commitment²³ to cover not only singular variables bound by the singular existential quantifier, but plural variables bound by the plural quantifier. I take it the primary motivation here is the feeling that it would somehow be unprincipled or arbitrary to refuse to extend the Quinean criterion to new quantifiers you introduce into your language: after all the same thing is going on in both cases; variables being bound by existential quantifiers.²⁴ However, this puts the cart before the horse as far as ontological commitment is concerned. If anything is going to be a primitive in ontology, it's going to be the notion of existence itself. Being is not going to be analysed away: to be is not *literally* to be the value of a variable. What I take to be attractive about the Quinean criterion is that it picks out all those entities that play some role in our fundamental theory and are therefore indispensable to it: if we felt the need to predicate something of an object, then it clearly plays a role in our theory. On the other hand, if we didn't predicate anything of it, it's not clear what role it could possibly play in the theory, and therefore why the theory should acknowledge its existence. The criterion does a nice job of counting up all the things we would intuitively suppose need to exist in order for our theory to be true.

However, things get more complicated when we bring plural quantification into the picture. Thinking for a moment about Yi's (1999) account of numbers as plurally instantiated properties might help to bring this out. Given plural properties, incurring a plural commitment to *some things that are (that is, have the property of being) seven* is – intuitively, at least – to incur a commitment to seven things, not to one thing, the 'some things' that are seven, or to eight things – the seven thing plus the 'some things'. This is supposed to be what distinguishes the plural strategy from various surrogate strategies that attribute the singular property of seven-ness to some further entity, in virtue of its relationship to the seven things, such as having them as parts or members. Seven-ness, on the plural view, is something that can only be instantiated jointly by seven entities. But if we're already committed to the seven things for independent reasons – for instance, because our theory says each of them possesses certain singular properties – then we shouldn't count them *again* (plurally this time, rather than individually) when we attribute to them, jointly, the property of being seven. The point is not that the plural quantifier isn't ontologically committing— I think it is. If you only had a plural quantifier in your ideology you could still read your ontological commitments off that quantifier alone: being ontologically committed to a plurality of seven things is not *less* ontologically committing than being ontological committed to seven things via singular quantification— so, strictly speaking it's wrong to say the plural quantifier is ontologically innocent. Rather, the point is that if you already

have the singular quantifier in your language, you don't incur extra ontological commitments just by adding another quantifier: both are quantifying over the same beings.

I think the really pressing challenge the nihilist needs to face is not that plural quantification complicates her ontology, but that it complicates her ideology: ideological commitments are incurred by adding primitive vocabulary to the language of your fundamental theory, and the plural existential quantifier needed by languages containing plural instantiation is just such a primitive— it is interdefinable with the plural universal quantifier, just as the singular existential quantifier is interdefinable with the singular universal quantifier, but is not definable in terms of the singular quantifiers (see Yi 2002).²⁵ If the nihilist has to expand her ideology in order to employ the Plural Instantiation Strategy, this cost might be weighed against the gains the nihilist claims for eliminating parthood from her ideology.²⁶

One way for the nihilist to respond is just to concede that there's no ideological parsimony gain to be had by swapping the parthood relation for plural quantification, but this doesn't diminish the importance of the ontological parsimony gained by eliminating composite objects from her ontology: belief in composition comes with both an ontological commitment to composites and an ideological commitment to parthood, whereas nihilism with plural quantification at least reduces ontological commitments – since, as I argued above, adding a plural quantifier to a language with singular quantification is ontologically innocent, even if the plural quantifier itself is ontologically committing. Now, I think this nihilist rightly points out an important gain in ontological parsimony, but is needlessly concessive concerning the ideological cost of plural quantification. The ideology objection all rests on what, exactly, the ideological cost of introducing the apparatus of plural quantification into your theory is, and I think the nihilist should say that the cost – if any – is negligible when you already have singular quantification in your ideology. The argument is structurally similar to the one I ran above concerning ontology: by having singular quantification in your ontology is already to incur the cost of having quantificational apparatus in your ontology, so although plural quantification is a significant ideological cost, it is not a significant *additional* cost for the nihilist relative to the believer in composition.

As Sider (2013) points out, and Goodman (1951) before him, not all ideological costs are created equal. However, if the nihilist wants to say that the cost of introducing a plural quantifier into her ideology is somehow minimal or negligible the burden of proof is on her. Such arguments can be tricky, given the many seemingly intractable issues surrounding theory choice methodology in general, but I think the intuitive case that can be made at least manages to shift the initial burden of proof away from the nihilist and onto her opponent. Let's say you had no quantificational apparatus in your theory at all: adding some quantificational apparatus into your theory would greatly enhance its power for a corresponding reduction in simplicity. However, let's say you're like Kris McDaniel (2010) or Jason Turner (2010), a fan of there being many ways of being— e.g. you think numbers exist in a different way from chairs (abstractly vs. concretely), or spacetime points in a different way from tropes (absolutely vs. relative to a spatial location), where this is not merely down to these things having different properties, but in literally existing in a different way. Turner and McDaniel both cash these ways of being talk out in terms of there being different fundamental quantifiers in one's ideology corresponding to each way that something could exist: thus, rather than having just the one fundamental quantifier, \exists , that ranges over everything, they have two quantifiers, \exists_C and \exists_A , say, that range over concrete and abstract things respectively (\exists can be defined in terms of \exists_C and \exists_A , but it is not – or does not have to be – part of their fundamental ideology). Adding a quantifier with a subscript shouldn't count as a significant addition to your ideology when you already have a quantifier in your theory: You already have quantificational apparatus in your language, you're just re-tasking that apparatus in a slightly different way. I would say it still *is* a complication to your ideology

compared to the person with just one quantifier, but it is not of the same order as adding in quantificational apparatus to a theory that lacks it entirely: By adding another quantifier when you already have one is not to *double* your ideological commitments related to quantification.

One way of further motivating this thought – at least in the particular cases under discussion – is that we’re multiplying tokens, not types. Many have followed Lewis²⁷ in thinking that, when it comes to ontology, qualitative parsimony matters, but quantitative parsimony does not— adding a new *type* of entity to your theory counts against it when judging its simplicity, while adding more entities of a kind already mentioned in your theory doesn’t make it any less parsimonious. One way in which you might apply the qualitative / quantitative distinction found in the discussion on ontological parsimony to ideology²⁸ is to say that adding a new bit of ideology to your theory is analogous to adding a new type of thing to your ontology, whereas the number of times you employ that ideology in your theory is analogous to adding a new thing of the same type to your ontology – e.g. the difference between adding a primitive modal operator or not, on the one hand, and saying of thousand and one things rather than a thousand things that they are possible, on the other. There may be some mileage in this analogy, but here’s another, perhaps equally useful one: adding a bit of ideology that is structurally different is adding something of a new type, whereas, adding a structurally identical bit of ideology is just adding more of the same rather than adding something of a new type. By structurally identical, I mean in the way – for instance – a distribution might be said to be structurally analogous. If I have three cakes and you have one, that distribution is structurally identical to the distribution in which you have three cakes and I have one: the form or type of distribution is just the same, the only difference is in what individuals fill the roles in that distribution. A new quantifier is just more quantifiers: it doesn’t matter whether the old one ranged over concreta but the new ranges over abstracta, or – indeed – whether the old one ranges over individuals and the new one ranges over pluralities.²⁹

So, is plural quantification as big a cost as having composition in your ideology?³⁰ As I’ve argued that the cost of introducing plural quantification into your ideology is negligible if you already have singular quantification in your ideology, the cost of having parthood in your ideology rather than plural quantification will be significant, unless the cost of adding parthood to your ideology is itself negligible. But there’s no reason to think that having parthood in your ideology is negligible – in fact, we have good reason to think that the cost is significant. The nihilist doesn’t have *any* mereological primitives in her language (Sider 2013), whereas the believer in composition needs at least one, and the jump from none to one is the most ideologically costly: it’s the difference between having no mereological structure in your theory at all and having a whole domain of mereological facts involving this primitive. It’s the equivalent of adding your first quantifier to your language to create a theory capable of expressing rich quantificational structure; or adding a primitive necessity or tense operator to your theory, in order to accommodate a domain of modal or temporal facts: once you have one primitive in a domain the expressive benefits of adding more are limited, and the ideological cost correspondingly more modest. If the believer in composition thinks their favoured mereological primitive is ideologically innocent it’s up to her to say why— and she can’t appeal to the same reasoning that the nihilist has available to show that plural quantification is no additional ideological cost.

To sum up, then, the plural instantiation strategy provides the nihilist with a way to answer the challenge from emergence that doesn’t rely on the empirically vulnerable premise that there is no such thing as emergence. By accommodating emergent properties in her ontology, this strategy also allows the nihilist to enjoy any of the theoretical benefits claimed by the believer in emergence. The cost of plural instantiation is its reliance on plural quantification: but I’ve argued that if you already have singular existential quantification in your language you don’t incur any extra commitments – ontological or ideological – by introducing plural quantification. So, in so

far as there is an ideological or ontological advantage to being a nihilist (which is not for us to adjudicate here), this is not eroded or cancelled by employing plural instantiation in order to account for emergent properties.³¹ In the next section I defend this view from a potential objection that the principle of ideological parsimony I appeal to undercuts any reason to think nihilism is parsimonious in the first place. In the final section, I want to discuss a worry that consciousness is especially resistant to plural instantiation and so will resist my general approach to dealing with emergent properties.

4. A Principle of Ideological Parsimony

How to decide when adding new ideology to your theory counts as a genuine theoretical cost or not? As explained in the previous section, it seems implausible to say that all additions to our ideology should be treated equally when evaluating a theory for ideological parsimony. The trouble is, it remains tempting to treat any given addition as just as costly as any other, since it's hard to come up with a principled and intelligible schema that weighs relative losses or gains of ideological parsimony in a more nuanced way. By arguing in the previous section that adding plural quantification to your ideology when you already have singular quantification is not a significant cost, this raises questions about the precise nature of, and justification for, the intended weighting schema I have obliquely appealed to. Saying more here is inevitably going to be tricky: I'll come clean right away that I don't have a precise and obviously correct method for weighing up ideological parsimony costs. What I propose instead, is to discuss a version of such a principle that has recently been proposed in the literature, saying what I think it gets right and, crucially, showing where it goes wrong. The discussion ensuing should, to some extent, better clarify and justify my own approach.

The view I have in mind is that of Sam Cowling (2013), which seeks to draw a distinction between *qualitative* ideological parsimony and *quantitative* ideological parsimony (analogous to the distinction often drawn between types of ontological parsimony). The basic idea is that complicating your theory with a new kind of ideology is costly, but adding more of the same kind of ideology is not: on the face of it, this is more or less the same line I was pushing in the previous section. However, if we're going to spell out the principle in terms of ideological kinds, then I believe Cowling to be far too liberal on what he counts as a kind. The stakes are high for me here as, if Cowling is correct about how a qualitative principle of ideological parsimony (qual-IP) would operate, then it turns out, according to Cowling, that nihilism itself is shown not have any ideological parsimony advantage over theories which have parthood in their ideology. Since it would be somewhat self-defeating of me to argue that the nihilist's best way to accommodate emergence does not lead to ideological bloat by appealing to a principle that would undercut any ideological parsimony advantages of nihilism from the outset, by arguing that Cowling's position on when we get ideology for free is too liberal I'm not just delineating my own position but responding to a potential objection to my entire strategy. Hence, a detailed discussion of Cowling's particular view is warranted.

Cowling initially motivates his case by appealing to the obvious innocence of equivalent or interdefinable bits of ideology: in the modal case, admitting a diamond operator into one's language is not an *extra* cost over and above a primitive box operator. So, *just counting* ideology in your theory isn't the way to go when weighing up ideological costs. Cowling's gloss on this: possibility is of the same ideological kind – the modal kind – as necessity, so once you have primitive modality of any sort all the rest comes for free. So far, this seems very much in keeping with sort of principle I want to appeal to in the previous section. However, Cowling claims a qualitative principle of ideological parsimony can be wielded to show that parthood is sufficiently similar to identity to be no cost over and above identity, thereby undercutting any

putative ideological parsimony gains for endorsing nihilism. Obviously, if such a principle is going to be so liberal as to count the cost of identity and parthood as one and the same, then any appeal to it will sabotage my attempt in the previous section to shore up nihilism's ideological parsimony advantage on the Plural Instantiation Strategy.

Can we make sense of the idea that composition is just like identity, only differing in the number or types of entities it concerns? Well, yes, easily enough, in that we can take composition to be the very same relation as identity, except holding between an individual and an (identical) plurality, rather than between an individual and an (identical) individual. But this is composition as identity in a strong form: the highly committal and controversial thesis that the whole is the very same thing as its parts. Sure, if we are prepared to assent to *that* then there's a strong claim that composition is no cost over and above identity: let's not quibble overmuch about the extra expressive power given to the identity relation here.³²

Yet Cowling seems to be presenting an argument that is neutral on strong CAI, and arguing from mere superficial similarities between parthood and identity that parthood is no cost over and above identity (2013, p.3906):

Like identity, it contributes nothing to the non-structural, qualitative character of the world, and, like identity, facts about its general nature seem to be a non-contingent matter. Furthermore, regardless of whether one endorses nihilism, classical extensional mereology demands certain conceptual ties between these relations. Most notably, the uniqueness of composition precludes distinct entities being composed of the very same objects.

While I'm perfectly happy to accept that parthood, along with identity, deserves the status of a 'broadly logical relation', and would even go further and assent to the claim that parthood is an identity-*like* relation in a meaningful sense, the gestured-to similarities are much looser than those present in the strong CAI case and, in my opinion, just not enough to support the claim that parthood is no cost over and above identity. Identity is a relation that never holds between distinct entities, while composition is a relation between some entities and some *further*, distinct entity – however closely qualitatively related the distinct entity may be to the other entities. To reject this fundamental difference between composition and identity is to endorse strong CAI. There's no sense here in which a well understood identity relation is simply being given enhanced expressive power to take more or different kinds of entities into its fold: rather, we have two entirely different, incompatible ideological primitives that are noted to be similar in interesting ways, but which must ultimately be treated separately. As such, the qualitative ideological parsimony principle as Cowling presents it does not provide a 'novel interpretation' (p.3906) of CAI in any substantive sense: either it should be read as a plea for full-fledged CAI along the usual lines, or as merely drawing attention to certain similarities between composition and identity which have no doubt contributed to the motivation behind weak forms of CAI in the past – either way, there is no new argument against nihilism.

So, is there any hope for extending a Cowling-esque principle of qualitative ideological parsimony beyond rather conservative observations about interdefinable operators: that no one wants to pay extra to employ a box operator in their theory if they're already committed to a diamond operator, and that debates concerning whether to *either* take the box operator *or* the diamond operator as primitive have the smell of non-substantivity. Cowling's proposal to use it to merge the ideological cost of two primitive notions that are in some sense similar yet ultimately independent of one another isn't going to fly: at least not without some further argument to back it up. However, I do think there's a plausible and interesting way, besides box and diamond type cases, in which an ideological notion can be no (significant) cost over and above another, closely related, ideological notion. This is where we start with an established ideological primitive and then augment our theory with a new notion that does basically the same

work as the old primitive but, say, involves objects from different ontological categories or involves a different number of objects or, more generally, allows that primitive to apply in different situations or to the same things in a different way. You may well then be in a situation to posit a new, more general, primitive that subsumes the original two – one that, from your new dialectical position can be cast as removing restrictions on the old primitive – or it may be that the old primitive can be defined in terms of the new more powerful or applicable primitive, such that it takes on the role of the general, subsuming primitive. In any case, this is obviously importantly different from the case of two *interdefinable* primitives and I take it we're on more controversial ground. This is because, despite the hope of being able to get rid of one primitive by defining it terms of the other being present in both cases, adding the new primitive increases the overall expressiveness of the theory, even if it ultimately subsumes, or is subsumed by, a different primitive of the theory.

Perhaps the best example to discuss, especially in the given context, is when you start out with a primitive that works in the manner of a one-one function and then add a primitive that does the same theoretical work, except in the manner of one-many function. For example: is it a complication of our theory to allow for variably polyadic relations? For instance, take the debate between Karen Bennett (2011) and L. A. Paul (2012), on the one hand, who argue broadly that composition or 'building relations' should be treated as a unified single relation across widely different theoretical roles and categories, despite holding between different numbers and kinds of *relata* in each case. On the other hand, Kris McDaniel (2010) argues that the composition relation should be considered a fundamentally different relation, with a different adicity, when employed for certain theoretical purposes compared to others: specifically, with spatial parthood holding between regions and sub-regions, while concrete parthood holds between a concrete whole, a part and a spatial region. The two sorts of views, while not straightforwardly incompatible, suggest two radically different pictures of the ideology of composition: one on which we have multiple primitives, each dealing with a fixed number of entities and entities in specific ontological categories; another on which we just have the one primitive that can deal with a varying number of entities across all relevant ontological categories. One thing which would look very bad to say is that we should simply dismiss the McDaniel picture out of hand, since going down this route will commit us to more than one parthood primitive. Here the quantitative/qualitative distinction could be put to work in explaining the intuition that whatever substantive issues may (or may not) be at stake between the Bennett-Paul picture and the McDaniel picture, that one does not have the advantage over the other on grounds of parsimony. Furthermore, though, by accepting that the posited composition relations in McDaniel belong to the same close-knit family of notions, we see that McDaniel doesn't have to pay a significant ideological cost for belief in both spatial composition *and* concrete composition compared to someone who *just* wanted to believe in spatial composition.

Plural Instantiation fits this mould: Here we start out with instantiation, which is standardly construed as holding one-one between an object and a property, and introduce plural instantiation, which is taken to relate a plurality of objects to a property in the same way that ordinary instantiation relates a single object to a property. To pick out the things that are to instantiate a given plural property, we similarly need to allow the existential quantifier to quantify not just over individual objects but also over pluralities: on the face of it, we complicate our ideology by adding a plural quantifier, but note that the old singular quantifier can be defined in terms of the new plural quantifier. Now, even on a *quantitative* conception of ideological parsimony, we get the result that there's no cost here, because of the interdefinability of the old primitive in terms of the new one— which suggests that the qualitative/quantitative distinction isn't quite getting to the heart of what's going on here. Nonetheless, from the preceding considerations we at least have a better grasp of the limits and potential applications of a principle of ideological parsimony that goes beyond the naïve 'counting primitives' approach.

5. *Fundamentality, Consciousness and Plural Instantiation*

We have seen that the plural instantiation strategy is a coherent and parsimonious way for the nihilist to overcome the objection from ontological emergence. In this final section, I want to make clear how these results advance the debate between nihilists and believers in composition and defend the strategy from an objection that Consciousness in particular cannot be plurally instantiated.

First up, it should be clarified exactly how the plural instantiation strategy for overcoming the objection from ontological emergence is different from the standard nihilist strategy for dealing with non-emergent phenomena of employing collective predicates to stand in for singular predications. According to the nihilist you are correct to say that there is chair in the room, despite chairs not existing, because ‘there are some simples arranged chair-wise’. This is not meant as a commitment to some new property of ‘being-arranged-chair-wise’, had collectively by the simples— a least, not a new fundamental property. Rather, we can truly predicate of the simples (collectively) that they are arranged chair-wise in virtue of the respective intrinsic properties of each member of the collection, taking into account the spatial relations (and other external relations if applicable) between each simple and every other: if there is a new property here in any sense, it is not something over and above how each simple is in itself and its spatial relation to each of the others. As such, we should say that there isn't any genuine plural instantiation going on at all: there is singular instantiation by simples and there are spatial relations between simples. Indeed, nihilists such as Sider (2013) and Cameron (2010*a*) will disavow the need to resort to paraphrases of the type ‘the simples are arranged chair-wise’, explicitly freeing them of the need include properties of being arranged F-wise or plural quantification in their fundamental theories.³³

Now, I've argued that plural quantification, and indeed plural instantiation, aren't a cost after all: so nihilists can keep the ideology and use it to accommodate genuine ontological emergence. This is my ultimate diagnosis of the worry that simples can't ‘co-operate’ to create consciousness even if simples can co-operate to, e.g., create light without composing some further thing:³⁴ if we're sure that any ‘higher-level’ property is fundamental, it's our own consciousness— right? But fundamentally, the thought goes, there are no plurally instantiated properties: so my consciousness must have a single bearer (me; my brain). However we've seen, firstly, that if ontologically emergent properties are possible, we have a reason to include the apparatus of plural instantiation in our fundamental ideology and secondly, that that apparatus isn't, in the final analysis, a significant cost.

Is recommending plural instantiation to the nihilist really to advance the debate? My strategy perhaps sounds similar to what some nihilists have had in mind all along as the solution to the problem of emergent consciousness. For instance, Dorr & Rosen (2002) write:

You don't exist; but the things you used to think of yourself as doing get done all the same. Certain atoms *jointly* think those thoughts, dream those dreams, and so forth. Things seem the way they do to those atoms jointly, but not to any single thing.

Such sentiments are all well and good. But they are also ambiguous between adopting the plural instantiation strategy, on the one hand, and thinking that such a strategy isn't needed because consciousness isn't emergent – that it is simply the matter of the correct arrangement of the intrinsic properties of individual simples and nothing over and above that – on the other. On the latter, as we've said, consciousness causes no more of a problem for the nihilist than that

walking is always a collective activity of many simples: there are some simples arranged walking-wise or arranged thinking-wise. This is, of course, a perfectly respectable position – but in this paper we have been interested in seeing whether nihilism can survive if we grant that there are genuinely ontologically emergent properties. If the former, then we are in agreement about the correct way to proceed, but merely gesturing to the possibility of collective activity does not by itself constitute a reply to the emergence objection, if plural instantiation turns out to be an incoherent or overly costly strategy – I hope what I have done in this paper is to show that employing plural instantiation is a *viable* strategy. I use Rosen and Dorr’s pronouncements only as an example: in general, there is a need for clarity on these issues – an appreciation of the various choice points – if we are to advance the debate.

One final objection I wish to meet head on concerns whether consciousness – which I’ve used throughout as a plausible example of a property that might be genuinely emergent – resists plural instantiation due to unique aspects of its own nature. If this is right, then there would be a worrying lack of general applicability in my approach for dealing with emergent properties—though it would still be applicable, as far as I can tell, for dealing with other plausible candidates, for example from Quantum Mechanics.

Is there something special about the nature of consciousness which means it requires a single bearer: a single subject?³⁵ My preferred diagnosis, above, is that this is down to nothing more specific about the nature of consciousness itself than the fact that emergent consciousness would be a fundamental property and so require the controversial apparatus of genuine plural instantiation. But perhaps it is rather that while genuinely plurally instantiated properties are perfectly possible, something very particular about the nature of consciousness makes it an exception to this general rule? The only viable candidate that springs to mind is the *unity* of consciousness: does the fact that all our experiences at a time seem to have a single, unified conscious subject, mean that consciousness itself needs to have a single bearer? While intuitive on first look, I find it to be a rather tenuous principle that experiential unity requires a literal unity – one thing – as a bearer, forcing conscious subjects to correspond one-one with objects in our ontology, rather than permitting our term ‘conscious subject’ to be a shorthand way to refer to a plurality of things involved in instantiating the same conscious property. Can we really introspectively tell that our consciousness is the product of one thing, one mereological whole, rather than some closely co-operating things?

The arbitrariness for the demand for what I shall call *merely mereological* unity – the need for the simples participating in an instance of consciousness to compose some further, single thing in order to successfully carry out this activity – can be brought out by seeing how demands for the requisite unity could, very plausibly, be either weaker or stronger. The weaker demand, which I need for my strategy to work, is that there be what I’ll call a high degree of *qualitative unity* involved in the bringing about of consciousness: they have similar intrinsic properties, perhaps, are spatially close together, are closely causally connected (and whatever else you need): the very sorts of conditions one might expect the restricted composition theorist to require in order for simples to compose a conscious composite object. The thing is, though, once these conditions are in place is there any sense in which the collection is made interestingly more unified by composing some further thing? After all, this further thing consists entirely of (albeit very close together, highly causally integrated) disparate, spatially separated objects each doing their own thing: qualitatively, there is no increase in unity. Excluding the property at issue – consciousness – we agree that the whole derives its properties from the parts, and so is a spatially discontinuous object composed of many disparate parts. The only difference is that when we point to the simples, according to the believer in composition, we are now also pointing to the thing the simples compose: but *qualitatively*, nothing has changed to make the conditions any more ripe for the instantiation of consciousness.

So, it really does come down to what criteria of unity you think is relevant here, and it's not at all clear how we should decide what's relevant. Nonetheless, you might suspect I am being perverse in pointing out that the change in the mereological facts doesn't precipitate any further change in qualitative facts towards greater unification - why should we expect it to? Well, there are plenty of other candidate bearers for consciousness which would seem to clearly beat wholes composed of atoms for inherent unity in both a qualitative *and* a mereological sense. Consider the Cartesian soul: mereologically speaking it is simple, partless— it is a pure unity, not even composed of a plurality of separate parts. Qualitatively, it is not a spatially discontinuous object: it is not even in space. Similarly, consider the view discussed briefly at the start of this paper - that the bearers of emergent consciousness could be extended emergent simples. Again, the extended simple is more unified than would be any mereological sum of the plurality that gave rise to it, in the sense that, being simple, it is metaphysically indivisible and cannot be partially one thing and partially another – partially my consciousness, perhaps, and partially a blancmange. The Mereological Unity Max criterion, as we'll call it, outdoes mere mereological unity at its own game. Plus, there is nothing stopping the extended simple from having greater qualitative unity: by not being spatially discontinuous, for example— so it's not absurd to demand that when some simples are arranged in the correct way to facilitate bringing about consciousness, some further qualitative change takes place in the world to enable this to happen. Third and finally, consider the view that the early emergentists set up their view in explicit opposition to: that to explain the non-mechanistic features of the world we need a singular entity of an entirely different kind from atoms and their sums – namely, entelechy.³⁶ Once again the idea is that a collection of mereological atoms just isn't a sufficiently unified or of the requisite nature to instantiate novel higher-level properties: and that the wholes they compose are hardly better placed to play this role either. Hence, if we're going by intuitions of requisite unity, the criteria could very plausibly be either stronger or weaker and it's hard to think of a principled reason for drawing the line precisely at the point that is favourable to traditional conceptions of ontological emergence but which causes trouble for the nihilist.

Emergentism is by design a middle way between mechanistic atomism and views which require some metaphysically special bearer for emergent properties, such as consciousness: fetishizing mereological sums as having the required unity to instantiate consciousness seems to me a bizarre hangover from the intuitions which motivated the positing of partless objects or objects of a new metaphysical category as bearers for non-mechanistic-seeming properties. Wholes are an uncontroversial addition to the emergentist's picture precisely because they are not seen as challenging the atomistic view of matter that emergentism takes as its starting point: they are no more mysterious than their parts. But this innocuousness which keeps them around in the emergentist picture also enables us to dispense with them if, as nihilists, we have independent reason – and the requisite theoretical apparatus – to exclude them from our ontology.

Whatever its initial ring of plausibility, the merely mereological criterion for the unity of conscious bearers is, I think, something that the nihilist would do well, on reflection, just to reject in the absence of some more compelling argument: just as the traditional emergentist eschewed long standing assumptions in favour of Mereological Unity Max. If you have a general systematic argument for eliminating most composite object (tables; stars; molecules) which leads you to endorse nihilism it's hard to see how shaky intuitions on the exact nature of the unity criteria should outweigh the promise of a more systematic theory. Not so, perhaps, if you are a restricted composition theorist who takes it to be constitutive of composition itself that it is precisely the sort of unity required for consciousness: that *what it is* for some things to compose some further thing is for them to compose a conscious subject. Here we are stipulating from the outset that the theoretical role of mereological composition is to pick out unified conscious subjects. But this would be a radical and unorthodox view that would be in need of some

independent argument: it is certainly not the view of van Inwagen, for instance, who is explicit that while being a whole is necessary for consciousness, being a conscious subject is not necessary for (let alone constitutive of) being a mereological whole.³⁷ Composition as Conscious Unity would be an interesting view which deserves attention in its own right, but it carries us far from the terms on which this debate is usually had and, certainly, beyond the scope and space constraints of this paper. As such, we should conclude that there is no good reason to think that the mereological unity involved in the composing of a single whole by its parts is sort of unity required for the instantiation of consciousness, emergent or otherwise.³⁸

Department of Philosophy
University of Leeds

NOTES

¹ The nature of the question itself and numerous possible answers are explored at length in van Inwagen (1990).

² For a defence of the latter view and discussion of what I have called the naïve answers see van Inwagen (1990). Other restrictivist answers to SCQ are developed by, e.g., Markosian (1998*b*), (2008) and Merricks (2001).

³ Defenders of universalism include: Armstrong (1997); Hudson (2001); Lewis (1986); Rea (1998); van Cleve (1986), (2008).

⁴ The view has been defended by, among others, Peter Unger (1979), (1980), Ross Cameron (2010*a*), Cian Dorr (2005), Dorr & Gideon Rosen (2002), Ted Sider (2013). Van Inwagen (1990) and Trenton Merricks (2001) put forward a strong general argument for endorsing nihilism but – as we we'll see – proceed to exploit lacunas in their arguments to rescue *living* or *conscious* objects (respectively) from elimination: as such, they are restrictivists about composition.

⁵ For instance, Peter Unger (1979), (1980) advances nihilism as a solution to the Problem of the Many, Sider (2001) recognises it as a competitor to Stage Theory for solving the paradoxes of coincidence, while Merricks (2001) is moved by the worry that wholes overdetermine the effects of their parts to present a general argument for the elimination of composite objects (that he finds an exception to the general argument for wholes that have powers over and above those of their parts is something we will move on to discuss shortly).

⁶ E.g. Cameron (2010*a*) defends a view on which nihilism can provide us with a minimal ontology; Sider (2013) argues for nihilism on the grounds that it is ideologically parsimonious. The latter will concern us in greater detail in §3.

⁷ See Timothy O'Connor & Hong Yu Wong (2012) for an overview of emergence. Broad (1925) is the most prominent of the early emergentists, notably preceded by Mill (1843).

⁸ Wilson (2015) provides a clear, exhaustive schema for understanding and characterising emergence, drawing on Kim's exclusion argument – see, e.g., Kim (1989) and (1999) – and building on many previous attempts in the literature. I will be assuming throughout this paper that ontological emergence should be understood as what Wilson characterises as *strong emergence*, remaining neutral on whether there is a putative problem for the nihilist of accommodating weak forms of emergence and whether, if so, my proposed solution is applicable to it.

⁹ See Barnes (2012) for a recent discussion of the nature of fundamental dependence and its importance in characterising ontological emergence.

¹⁰ McLaughlin (1992) argues that the decline of British Emergentism was due to quantum theory providing a complete explanation for 'higher level' phenomena in chemistry and biology, thereby undercutting a motivation for positing ontologically emergent chemical or biological properties or laws. However, while quantum theory enables us to see how, e.g., chemistry is governed by the same dynamics as fundamental physics, these generally applicable dynamics treat multi-particle systems in a way that suggests their properties are not fully determined by their subsystems. The resulting picture of emergence is one that eschews the traditional picture of special science laws or properties emerging from the

fundamental laws or properties in favour of one where emergence appears to be somehow imbedded in the structure of the fundamental laws or properties themselves.

¹¹ Compare Maudlin (2007, pp.50-77). For more detailed discussion of the phenomena see, e.g., Esfeld (2001) or Penrose (2004).

¹² The important thing here is that van Inwagen finds at least one property of the whole not derived from the property of the parts. I discuss in §5 whether van Inwagen takes the nihilist not to be able to accommodate emergent properties generally, or just consciousness in particular.

¹³ Schaffer leaves open that ‘entanglement relations’ could be posited – new external relations between the simples specifically posited to ground their collective spin – but this makes the assumption that what we took to be fundamental monadic spin properties of wholes can be explained away by stipulating that the intrinsic properties of the parts plus a brand new external relation can play the equivalent theoretical role. I therefore count this as explaining away emergence rather than accommodating it (it’s fine if you think you can explain away emergence, but here we are narrowing our focus to whether it can be accommodated).

¹⁴ See Schaffer (2010) and Barnes (2014) respectively.

¹⁵ See Markosian (1998*a*) for an in depth discussion of the nature of simples – especially whether they are necessarily point-sized.

¹⁶ See Parsons (2004); McDaniel (2007); Cameron (2010*b*).

¹⁷ An interesting version of something in the vicinity of this strategy has recently been defended by Cotnoir (2013), relying on the possibility of extended simples and claiming to allow the nihilist to accommodate emergent properties. However, his parts as counterparts approach requires some explicit heavy duty commitments to both substantial space and fundamental distribution properties (p.237). Although there is an interesting discussion to be had on the relative merits of my strategy and Cotnoir’s, I have decided to use this paper to develop my own view fully rather than taking a comparative approach. If my view is viable it is at the very least an interesting result that a strategy exists that can do without Cotnoir’s not insignificant theoretical commitments (that’s not to say my view is without its own commitments, but merely that they are quite different—opening up the range of options for the nihilist).

¹⁸ This is assuming that she believes in a world that contains simples rather than one that is ‘gunky’.

¹⁹ In drawing this distinction I follow Yi (2002). Elsewhere in the literature the term ‘plural instantiation’ is sometimes used to refer to multiple instantiation.

²⁰ E.g. Morton (1975); Mundy (1989); Rayo (MS).

²¹ Byeong-Uk Yi provides arguments for adopting plural instantiation which primarily concern its ability to give an attractive metaphysical account of numbers.

²² For a long list of doubters see Linnebo (2004).

²³ See Quine (1948).

²⁴ Compare Boolos (1984).

²⁵ An alternative way to look at the extra cost involved here is that the plural quantifier is more internally complex than the singular quantifier: as well as having a primitive logical predicate of *identity* the believer in plural quantification needs the primitive predicate *is one of*, to denote that a given object belongs to a plurality. So, the real cost of plural quantification is adding more identity-like structure to the theory (the ‘is one of’ relation), rather than more quantificational structure as such (for once one has the more expressive plural quantifier arguably one does not need the singular quantifier too). However, I think what follows will play out much the same regardless of your view you on the technicalities of where precisely the extra commitment lies.

²⁶ Sider (2013) argues for the view that eliminating parthood is an ideological economy.

²⁷ e.g. Lewis (1973, p.87) – though compare Nolan (1997).

²⁸ A view I’ve come across a number of times in conversation.

²⁹ A very different way of arguing for this conclusion – which I’m not recommending, but is worth drawing attention to – is to motivate a retreat from the very austere logical realism of Sider (2011). One could still think that many ideological commitments are a theoretical cost – such a primitive modal or tense operators, say – but hold the more general background features of our logical, such as identity and quantification, are not *wordly* features as such, but just necessary artefacts of theorising or thought, and therefore innocent.

³⁰ I'm assuming her that 'compose' is not a primitive in the believer's theory – since composition itself a plural relation of the parts to the whole – and that it can be explained in terms of a singular relation like parthood.

³¹ In this case, at least, we are not simply 'pushing a bump around under the carpet' as Bennett (2009, p.65) suggests the nihilist is prone to do— rather, expressive power is gained without a corresponding loss of parsimony.

³² See Baxter (1988*a*) and (1988*b*) for canonical formulation of CAI. See, e.g., Wallace (*draft*), (2011*a*), (2011*b*) for recent discussion and defence.

³³ Thus, says Sider (2013), exempting this sort of nihilist from the bloat of ideology and properties associated with need to provide paraphrases identified by Bennet (2009). It's important to emphasise that even if you are a nihilist who thinks that paraphrase apparatus, such as plural quantification, is heavyweight ideology you are genuinely committed to this does not make your paraphrases equivalent to plural instantiation, for the reasons just outlined: something being 'arranged F-wise' is a derivative property nothing over and above how the things in question are arranged.

³⁴ See van Inwagen (1990, p.188).

³⁵ Thanks to an anonymous referee for pressing me to address this issue directly.

³⁶ See McLaughlin (1992) for discussion of how Broad (1925) endorsed emergence as a middle way between mechanist views and substance vitalist views such as Driesch's (1914).

³⁷ Van Inwagen (1990: p.120).

³⁸ Thanks to Elizabeth Barnes, Ross Cameron, Christopher Daly, Scott Shalkowski and an anonymous referee for invaluable comments on earlier drafts of this paper, as well as to audiences at the University of Leeds Centre for Metaphysics and Mind Seminar and the University of Leeds PG Seminar for immensely useful discussion. Much of the work towards this article was enabled by the support of the Arts and Humanities Research Council.

REFERENCES

- Armstrong, D. M. (1997). *A World of States of Affairs* (New York: Cambridge).
- Barnes, Elizabeth (2012). 'Emergence and Fundamentality', *Mind*, 121 (484): 873-901.
- Barnes, Elizabeth (2014). 'Fundamental Indeterminacy', *Analytic Philosophy*, 55 (4): 339–362.
- Baxter, D. L. M., (1988*a*). 'Identity in the Loose and Popular Sense', *Mind*, 97: 576–582.
- Baxter, D. L. M., (1988*b*). 'Many-One Identity', *Philosophical Papers*, 17: 193–216.
- Bennett, Karen (2009). 'Composition, colocation, and metaontology', in David John Chalmers, David Manley & Ryan Wasserman, eds., *Metametaphysics: New Essays on the Foundations of Ontology* (New York: Oxford University Press), pp.38-76.
- Bennett, Karen (2011). 'Construction area (no hard hat required)', *Philosophical Studies*, 154: 79–104.
- Boolos, George (1984). 'To be is to be a value of a variable (or to be some values of some variables)', *Journal of Philosophy*, 81 (8): 430-449.
- Broad, C. D. (1925). *The Mind and Its Place in Nature* (London: Routledge & Kegan Paul).
- Cameron, Ross (2010a). 'How to have a radically minimal ontology', *Philosophical Studies*, 151: 249-64.

- Cameron, Ross (2010*b*). 'Truthmaking for Presentists', in Karen Bennett and Dean W. Zimmerman, eds., *Oxford Studies in Metaphysics*, 6 (New York: Oxford University Press).
- Cotnoir, A. J. (2013). 'Parts as counterparts', *Thought*, 2: 228–241.
- Cowling, Sam (2013). 'Ideological Parsimony', *Synthese*, 190: 889-908.
- Dorr, Cian (2005). 'What We Disagree About When We Disagree About Ontology', in M. Kalderon, ed., *Fictionalism in Metaphysics* (Oxford: Oxford University Press), pp.234–86.
- Dorr, Cian & Rosen, Gideon (2002). 'Composition as a Fiction', in Richard Gale, ed., *The Blackwell Companion to Metaphysics* (Oxford: Blackwell), pp.151-174.
- Driesch, Hans (1914). *Problem of Individuality* (London: Macmillan).
- Esfeld, Michael (2001). *Holism in Philosophy of Mind and Philosophy of Physics* (Dordrecht: Kluwer).
- Goodman, Nelson (1951). *The Structure of Appearance* (Cambridge: Harvard University Press).
- Hudson, H. (2001). *A Materialist Metaphysics of the Human Person* (London: Cornell UP).
- Kim, Jaegwon (1989). 'The Myth of Nonreductive Materialism', *APA Proceedings*, 63: 31-47.
- Kim, Jaegwon (1999). 'Making Sense of Emergence', *Philosophical Studies*, 95 (1): 3-36.
- Lewis, David (1973). *Counterfactuals* (Oxford: Basil Blackwell).
- Lewis, David (1986). *Philosophical Papers*, ii (Oxford: Oxford University Press).
- Linnebo, Øystein (2004). 'Plural Quantification', in Edward N. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Spring 2013 edition). URL = <http://plato.stanford.edu/archives/spr2013/entries/plural-quant/>
- Markosian, Ned (1998a). 'Simples', *Australian Journal of Philosophy*, 76: 213-26.
- Markosian, Ned (1998b). 'Brutal Composition', *Philosophical Studies*, 92 (3): 211-249.
- Markosian, Ned (2008). 'Restricted Composition', in J. Hawthorne, T. Sider, & D. Zimmerman, eds., *Contemporary Debates in Metaphysics* (Malden: Blackwell), pp.341-364.
- Maudlin, Tim (2007). *The Metaphysics within Physics* (New York: Oxford).
- Merricks, Trenton (2001). *Objects and Persons* (New York: Oxford).
- McDaniel, Kris (2007). 'Extended Simples'. *Philosophical Studies*, 133: 131-141.
- McDaniel, Kris (2010). 'A Return to the Analogy of Being', *Philosophy and Phenomenological Research* LXXXI (3): 688 – 717.
- McLaughlin, Brian (1992). 'The Rise and Fall of British Emergentism,' in A. Beckermann, H. Flohr and J. Kim, eds., *Emergence or Reduction?* (Berlin: Walter de Gruyter).

- Mill, J. S. (1843). *System of Logic*, (London: Longmans, Green, Reader, and Dyer) [8th ed., 1872].
- Morton, Adam (1975). 'Complex Individuals and Multigrade Relations.' *Nous*, 9 (3): 309-18.
- Mundy, Brent (1989). 'Elementary Categorical Logic, Predicates of Variable Degree, and a Theory of Quantity', *Journal of Philosophical Logic*, 18: 115-140.
- Nolan, Daniel (1997). 'Quantitative Parsimony', *The British Journal for the Philosophy of Science*, 48: 329-343.
- O'Connor, Timothy & Wong, Hong Yu (2012). 'Emergent Properties', in Edward N. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Spring 2012 edition). URL = <http://plato.stanford.edu/archives/spr2012/entries/properties-emergent/>
- Parsons, Josh (2004). 'Distributional Properties' in F. Jackson and G. Priest, eds., *Lewisian Themes*. (Oxford: Oxford University Press), pp.173-180.
- Paul, L. A. (2012). 'Building the world from its fundamental constituents', *Philosophical Studies*, 158: 221–256.
- Penrose, Roger (2004). *The Road to Reality* (London: Vintage).
- Quine, W. V. (1948). 'On What There Is', *Review of Metaphysics*, 2 (5): 21-36.
- Rayo, Agustin (MS). *Plural Predication*.
- Rea, M. (1998). 'In Defence of Mereological Universalism', *Philosophy and Phenomenological Research*, 58 (2): 347-360.
- Schaffer, J. (2007). 'From Nihilism to Monism', *Australasian Journal of Philosophy*, 85 (2): 175-191.
- Schaffer, Jonathan (2010). 'Monism: The Priority of the Whole', *Philosophical Review* 119 (1): 31-76.
- Sider, Theodore (2001). *Four Dimensionalism: An Ontology of Persistence and Time* (New York: Oxford University Press).
- Sider, Theodore (2011). *Writing the Book of the World* (New York: Oxford University Press).
- Sider, Theodore (2013). 'Against Parthood', in Karen Bennett and Dean W. Zimmerman, eds., *Oxford Studies in Metaphysics* 8 (New York: Oxford University Press).
- Turner, Jason (2010). 'Ontological Pluralism', *The Journal of Philosophy* 107 (1): 5-34.
- Unger, P. (1979). 'There are no Ordinary Things,' *Synthese*, 41 :117–154.
- Unger, P. (1980). 'The Problem of the Many', *Midwest Studies in Philosophy*, 5 (1): 411-468.
- Van Cleve, J. (1986). 'Mereological Essentialism, Mereological Conjunctivism, and Identity Through Time', *Midwest Studies in Philosophy*, 11 (1): 141-156.

Van Cleve, J. (2008). 'The Moon and Sixpence: A Defence of Mereological Universalism', in J. Hawthorne, T. Sider and D. Zimmerman, eds., *Contemporary Debates in Metaphysics* (Oxford: Blackwell), pp.321-340.

Van Inwagen, Peter (1990). *Material Beings* (New York: Cornell University Press).

Wallace, Meg (2011*a*). 'Composition as Identity: Part 1', *Philosophy Compass*, 6 (11): 804-816.

Wallace, Meg (2011*b*). 'Composition as Identity: Part 2', *Philosophy Compass*, 6 (11): 817-827.

Wallace, Meg (*draft*). 'The Many Advantages of Composition as Identity'.

Wilson, Jessica (2015). 'Metaphysical Emergence: Weak and Strong', in Bigaj, Tomasz and Wuthrich, Christian, eds., *Metaphysics in Contemporary Physics: Poznan Studies in the Philosophy of the Sciences and the Humanities* (Amsterdam: Rodopi).

Yi, Byeong-uk (1999). "Is Two a Property?", *Journal of Philosophy*, 96:163-90.

Yi, Byeong-uk (2002). *Understanding the Many* (New York: Routledge).

Yi, Byeong-uk (2005). 'The Logic and Meaning of Plurals. Part I', *Journal of Phil Logic*, 34: 459-506.

Yi, Byeong-uk (2006). 'The Logic and Meaning of Plurals. Part II', *Journal of Phil Logic*, 35: 239-88.