

This is a repository copy of *Manyism as mereologicism*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/229261/>

Version: Accepted Version

Article:

Trueman, Rob orcid.org/0000-0002-2908-7985 (2025) Manyism as mereologicism. Synthese. 99. ISSN: 0023-7857

<https://doi.org/10.1007/s11229-025-05123-2>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

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.

Manyism as Mereologicism

Robert Trueman & Simon Thunder

There is a widespread intuition that mereology should be *ontologically innocent*: the existence of a fusion should be nothing over and above the existence of its parts; if you are already committed to some parts, their fusion should not count as an extra ontological commitment. But this intuition is problematic. A (proper) fusion is distinct from each of its (proper) parts,¹ so how can its existence involve nothing more than the existence of those parts? Surely it involves the existence of *one more thing*!

In this paper, we will compare two attempts to solve this problem. Both of them start from the same place. They both try to make mereology ontologically innocent by identifying a fusion with its parts: although the fusion is distinct from *each* of its parts, it is identical to *all* of them taken together.² But they disagree over the logical status of fusions. According to *Composition as Identity* (CAI), a fusion is a genuine *individual*, despite its being identical to a *plurality* of parts;³ in this sense, a fusion is both *one* and *many*.⁴ According to *manyism* (Thunder 2023), a fusion is just a plurality, not an individual as well; in this sense, a fusion is *merely many*.

We will argue that CAI cannot deliver the ontological innocence of mereology without violating some independently attractive desiderata. But manyism can. In fact, when the details are appropriately filled out, manyism yields (what we will call) *mereologicism*, the result that all of the axioms of Classical Mereology are logical theorems. Given manyism, mereology is innocent because logic is.

We will begin by clarifying how manyism differs from CAI in §1. Then, in §2, we will argue that the version of CAI presented in §1 cannot make mereology ontologically innocent without imposing an unacceptable restriction on the principle of Plural Comprehension. Now, if we were to end the paper there, our conclusion would be a little stale, since a number of authors have already argued that combining CAI with unrestricted Plural Comprehension has disastrous consequences.⁵

¹ A fusion is *proper* iff it has a proper part (i.e. a part other than itself). In this paper, we will always use ‘fusion’ and ‘part’ to mean *proper fusion* and *proper part*, unless we clearly indicate otherwise.

² This is not the only possible approach. You might instead try to solve the problem by saying that fusions are *grounded* in their parts (see e.g. Cameron 2014). Officially, this kind of solution simply falls outside the remit of this paper. *Unofficially*, we worry that it is not really a solution, but just a label for our problem. Grounding is meant to underwrite metaphysical explanation. However, to say that the existence of a fusion is grounded in the existence of its parts is not actually to give an explanation; it is just to promise that there is one to be given. What we want to know is precisely *how* the existence of one thing — a fusion — could be grounded in the mere existence of some other things — its parts. (Also see Barker (2021) for reasons to be suspicious of the whole idea that grounded entities are ontologically innocent relative to their full grounds.)

³ Throughout this paper, we will use ‘plurality’ and its cognates as *pseudo-singular* devices, to allow easy discussion of CAI and manyism in English: talking about a ‘plurality of things’ is just a helpful, syntactically singular, way of talking plurally about the things themselves.

⁴ At least, this is how CAI is usually understood. We are happy to allow that, in a broader sense, the manyism described below could also be seen as an alternative version of CAI.

⁵ See Yi 1999; Sider 2007: §3, 2014; Calosi 2016; Loss 2018.

However, as we will show in §3, our argument can be reworked against newer versions of CAI that are immune to the familiar objections: in §3.1, we will argue that Loss’s (2021a,b, 2022) *Atomic Composition as Identity* is vulnerable to our argument more-or-less as formulated in §2; and in §3.2, we will argue that Cotnoir’s (2013) *Composition as General Identity* cannot make mereology ontologically innocent without restricting Leibniz’s Law. In fact, our argument can be generalised to show that *no* version of CAI can make mereology ontologically innocent without restricting Plural Comprehension or Leibniz’s Law, as we will show in §3.3. Finally, in §4, we will argue that manyism can deliver the ontological innocence of mereology without restricting either principle. The key to our argument will be demonstrating that we can use manyism to establish mereologicism.

1 CAI versus manyism

Consider your favourite mereologically complex mug. CAI and manyism have this much in common: they both identify the mug with its parts. (Actually, manyism as we understand it here only identifies the mug with its *atomic* parts, but that is a difference that we don’t need to dwell on right now.) So CAI and manyism agree that there is an important sense in which the mug is *many* things: its (atomic) parts. However, they disagree about whether there is also a sense in which the mug is *one* thing, an individual: CAI says that there is, and manyism says that there isn’t.

1.1 The logical difference

It can be difficult to make sense of CAI, and even more difficult to see how it differs from manyism. We can start to make things clearer by thinking about how to formalise CAI. We will need to use a formal system that allows us to express identities between individuals and pluralities.⁶ The natural suggestion is a single-sorted plural logic, in which one sort of first-order variable (e.g. x, y, z) can take individuals and pluralities as values. (This is, in effect, what Sider (2007: 57) recommends.)⁷ It is important to acknowledge that the choice to use a single-sorted system is not philosophically neutral. In such a system, whatever can be said of an individual can also be said of a plurality, and *vice versa*; but you might reasonably doubt that this is always possible. However, we do not have any concerns about this expressive freedom,⁸ and crucially, it seems to be part and parcel of the CAI picture.

⁶ Van Inwagen (1994) was sceptical about the very intelligibility of identifying an individual with a plurality. However, we are willing to grant, at least for the sake of argument, that we can make sense of such identifications; for us, the question is just how *best* to make sense of them.

⁷ CAI-ists often *appear* to use two sorts of variables (e.g. single-letters for individuals — x, y, z — and double-letters for pluralities — xx, yy, zz). However, if we can express identities between pluralities and individuals, then Leibniz’s Law (as formalised below) requires that whatever can be said of an individual can also be said of a plurality, and *vice versa*. At this point, there can be no objection to introducing an *unsorted* variable, which may be syntactically substituted for either sort of variable, and which takes both individuals and pluralities as its values. (For related discussion, see Button and Trueman 2022: §4.)

⁸ For further discussion, see: Florio and Linnebo 2021: §11.7; Button and Trueman 2022: §4.2.

In a single-sorted setting, we do not distinguish between individuals and pluralities with different sorts of variables, but with two predicates, I and P . We can then offer the following initial formalisation of CAI:

$$\text{CAI: } \forall x \forall y (Fxy \leftrightarrow (Ix \wedge x = y)),$$

where Fxy symbolises x is a fusion of y , including all improper cases.⁹ As we will see in §3, some CAI-ists have tweaked this formalisation in various ways, but it is enough to get us started.¹⁰

For now we will assume that numerical identity obeys *Leibniz's Law* (although we will discuss that a little further in §3.2):

Leibniz's Law: $\forall x \forall y (x = y \rightarrow (\phi \leftrightarrow \psi))$, whenever ψ is a result of substituting a free occurrence of y for a free occurrence of x in ϕ .¹¹

Given Leibniz's Law, CAI implies that the fusion of some parts is itself both an individual and a plurality. Some critics of CAI (e.g. Yi 2014) have complained that this is a straightforward contradiction, because they analyse Px and Ix as x includes at least two things and x does not include at least two things, respectively. However, CAI-ists should just reject these analyses. What really matters is how I and P interact: pluralities are precisely pluralities of individuals.¹² CAI-ists are free, then, to leave I as primitive, and then define P as follows:

$$\forall x (Px \leftrightarrow (\exists y \exists z (y \neq z \wedge y \varepsilon x \wedge z \varepsilon x) \wedge \forall y (y \varepsilon x \rightarrow Iy))),$$

where $y \varepsilon x$ symbolises y is vertically included in x . (Vertical inclusion should be contrasted with *horizontal* inclusion: vertical inclusion is the relation that a plurality bears to each of the things that it groups together; horizontal inclusion is the relation that a plurality bears to each of its subpluralities. Throughout this paper, we will use 'inclusion' and its cognates for vertical inclusion.¹³) In other words, I is just our label for the things that are grouped together in pluralities; and P is just our label for the things that group together individuals.¹⁴

⁹ Improper cases are included in the following sense: if Ix then Fxx , even if x has no proper parts.

¹⁰ Our CAI is what Sider (2007: 59) calls *superstrong* CAI. Merely *strong* CAI consists in the left-to-right direction of the biconditional only. However, as Sider (2007: 60 fn. 24) notes, the right-to-left direction follows from Leibniz's Law and the fact that every individual is a (possibly improper) fusion of itself. We may actually want something stronger than our CAI — e.g. a strict biconditional or a higher-order identity — but we will stick with a material biconditional just to keep things simple.

¹¹ As is standard, we assume that indirect contexts, like 'Sharon believes that...' and 'Joseph hopes that...', are exempted from this scheme. (So when we consider restricting Leibniz's Law in §3.2, we mean restrictions that go beyond these standard exemptions.)

¹² This is only really true of first-level pluralities, not the super-pluralities that we will introduce in §4. However, we will only use P for first-level pluralities.

¹³ The distinction between vertical and horizontal inclusion is due to Oliver and Smiley (2016: ch.15); as they explain, vertical inclusion is like set-membership, and horizontal inclusion is like subthood. The vertical/horizontal distinction is not always drawn by plural logicians. It is instead common to work with a single notion of inclusion which, from our point of view, is the disjunction of vertical and horizontal inclusion. However, it is important that CAI-ists not slur this distinction: by identifying pluralities with individuals, they make it possible for one plurality to be vertically included in another. (The vertical/horizontal distinction will also be essential to our formalisation of manyism in §4.2.)

¹⁴ Loss (2021b: §7) and Payton (2021b) offer similar responses to Yi, albeit formulated in (superficially — see fn. 7) two-sorted settings.

So, according to CAI, the mereological fusion of some parts appears at two places in the plural structure: it is both a plurality and also an individual, the sort of thing that can itself be grouped into a plurality. Understood in this way, CAI may be consistent, but it is certainly still non-standard. Ordinarily, pluralities are assumed to be distinct from individuals.¹⁵ In this regard, then, CAI is logically revisionary. By contrast, manyism is logically unexceptional: according to manyism, fusions appear *only* at the level of pluralities; they are *not* also found at the level of individuals.

Later, in §4, we will develop manyism in much more detail. However, this initial sketch is already enough for us to draw an important philosophical distinction between manyism and CAI. It turns out that, despite their similarities, these views are motivated by two importantly different perspectives on ontology.

1.2 The ontological difference

What would ever motivate a philosopher to prefer CAI, with its non-standard approach to plural logic, over the much more straightforward manyism? Well, it can be very hard to hear the manyist's denial that fusions are individuals as anything other than the denial that they exist at all. In other words, it can be hard to see how manyism is anything other than a souped-up form of nihilism. Here is how Cameron puts it, after describing a view which comes very close to our manyism:

[Manyists] might talk the CAI-ist talk, but it's clear that they have the mereological nihilist's ontology, and that everything else is just a way of talking. But CAI is not meant to be mereological nihilism with some fancy talk, it's meant to be a radical ontological thesis. It seems to me that what distinguishes CAI as an interesting ontological thesis from the above nihilist theory is their additional claim that not only is everything identical to itself and that every collection of things is identical to that collection of things, but that there are collections of things identical to some *one* thing. That is, it is the postulation of many-*one* identity that makes CAI interesting — that distinguishes it from mereological nihilism with some fancy talk. (Cameron 2012: 547)

This is the only good reason we can offer for choosing CAI over manyism: you think that manyism collapses into nihilism, and you don't want to be a nihilist.¹⁶

Now, maybe there's nothing wrong with being a nihilist, especially one who knows how to talk fancy. But, for the record, we manyists do not think of ourselves as nihilists. We say that mereological fusions exist, and we mean that just as seriously as anyone else. It's just that we don't think that fusions are individuals.

If that strikes you as giving with one hand and taking with the other, then that is because you are working with a *demanding* conception of ontology, according to which only individuals exist: only quantification over individuals is ontologically

¹⁵ More precisely: *proper* pluralities (i.e. pluralities that include more than one individual) are usually assumed to be distinct from individuals. Some plural systems identify each individual with its 'singleton plurality'. However, in this paper, all pluralities are proper pluralities.

¹⁶ For a review of some other reasons a philosopher might give for preferring CAI, including explanations of why they are not *good* reasons, see Thunder (2023). It is also worth noting that, given some plausible background assumptions, CAI itself implies nihilism, at least when it is formulated as above (Calosi 2016; Loss 2018; Yi 2021); however, CAI-ist can avoid this unwanted result by rejecting those background assumptions, or by revising their formulation of CAI (as we discuss in §3).

committing; we are not ontologically committed to any plurality we quantify over, unless we take the further step of identifying it with an individual. But manyists have a *liberal* conception, according to which pluralities exist in exactly the same sense as individuals: we are automatically ontologically committed to any plurality we quantify over, just as we are to any individual.¹⁷

The distinction between these conceptions can be obscured by the fact that it is often natural to formulate existential claims in plural terms. Regardless of which conception you prefer, you might say, for example, that electrons exist, or even that a *plurality* of electrons exist. However, as ordinarily understood, that is just a way of committing yourself to the existence of each individual electron in the plurality. Our concern here is with pluralities *themselves*: are we ontologically committed to the plurality of electrons itself, alongside each individual electron? Now that it has been clarified, this might sound like a bad question. Ever since Boolos (1984) popularised plural logic, one of its main selling points has supposed to be that it gives us a way of collecting together individuals without introducing *collections* as things in their own right. However, in order to make sense of CAI, we have adopted a single-sorted plural logic. In this setting, whatever can be said of an individual can be said of a plurality too. So it makes sense to ask whether pluralities themselves exist, and our two conceptions of ontology answer this question in different ways. According to the demanding conception, only individuals exist; so a plurality exists *only if* it is identical to an individual. According to the liberal conception, on the other hand, every plurality automatically exists, whether or not it is identical to an individual.

We can think of both conceptions as descending from Quine's (1948) famous criterion of ontological commitment: *to be is to be a value of a variable*.¹⁸ When Quine presented this criterion, he intended us to apply it in the context of a singular first-order logic. In that setting, then, Quine's criterion tells us that a theory's ontology includes all and only the individuals that it quantifies over. Everything else (e.g. the way that the theory categorises individuals with predicates) belongs to the theory's *ideology*, i.e. the suite of resources it makes available for describing the individuals in its ontology.

However, we are now working in a single-sorted plural logic. In this system, we have just one sort of variable, which indiscriminately takes individuals and pluralities as values. Now that we have made this change, we are faced with a choice about how to apply Quine's criterion in our new setting.¹⁹

We have two options. We can either preserve the ontology that Quine demarcated, or the way in which he demarcated it. To choose the first option is to adopt the demanding conception of ontology, according to which only individuals exist. So existence is not really expressed by the unrestricted existential quantifier: only something of the form $\exists x(Ix \wedge \phi)$ expresses ontological commitment; $\exists x(Px \wedge \phi)$ expresses some other sort of commitment — call it an *ideological commitment* — about

¹⁷ See Florio and Linnebo (2016, 2021: ch. 8) for a related discussion of 'broad' and 'narrow' ontological commitments.

¹⁸ For related discussion of manyism and Quine's criterion, see Thunder (2023: §2.2.2).

¹⁹ Of course, Quine himself would have had no truck with a single-sorted plural logic. So this is very much a choice for *lapsed* Quineans.

how the individuals in our ontology can be grouped together.²⁰ For this reason, we should be wary of talking about pluralities ‘existing’ when we’re working with the demanding conception; that is really appropriate only if those pluralities also happen to be individuals. However, we find that terminological policy unsustainable, and so we’ll just have to remember that, on this conception, we’re not meant to take talk of pluralities’ existing too seriously.

If we instead choose to stick with the *way* in which Quine demarcated a theory’s ontology, then we are led to the liberal conception: the values of our variables include pluralities as well as individuals, and so our ontology must include them too. It is important to stress that, although this option yields a decidedly non-Quinean ontology, it still retains a great deal of what mattered to Quine. For example: existence is still univocal, and it is still expressed by the existential quantifier, $\exists x\phi$. For another example: we can still make identifications and draw distinctions throughout our ontology, since the variables in $x=y$ can each take individuals and pluralities as values.²¹

To be clear, our aim is not to force anyone to adopt our liberal conception of ontology. Indeed, we are not sure that anything could *force* that: it seems to us that we are simply confronted with a choice about how best to extend Quine’s criterion to deal with a new logical setting. Instead, our point is just that a philosopher should prefer CAI to manyism only if they adopt the demanding conception (and don’t want to be a nihilist). So, in what follows, we will take the demanding conception of ontology to be folded into CAI, and likewise the liberal conception to be folded into manyism.

2 Why CAI cannot deliver ontological innocence

Here are some very well known remarks that Lewis made in favour of CAI:

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 of their fusion. The new commitment is redundant, given the old one [...]

²⁰ Those who adopt the demanding conception of ontology may prefer not to use the term ‘ideological’ to characterise their commitment to pluralities. For example, they might prefer to follow Rayo (2007) and Payton (2022) in calling them ‘plethological’ commitments. Nothing turns on the choice of terminology here: all that matters is that, according to the demanding conception, $\exists x(Px \wedge \phi)$ does not yet express an ontological commitment.

²¹ It is worth noting that we can only make sense of a substantial disagreement between the demanding and liberal conceptions of ontology relative to a *single-sorted* plural logic. Contrast Trueman’s (2021: ch. 7) discussion of the ontological significance of different types of quantification.

I say that composition — the relation of part to whole, or, better, the many-one relation of many parts to their fusion — is like identity. The ‘are’ of composition is, so to speak, the plural form of the ‘is’ of identity. Call this the Thesis of *Composition as Identity*. It is in virtue of this thesis that mereology is ontologically innocent: it commits us only to things that are identical, so to speak, to what we were committed to before. (Lewis 1991: 81–2)

Lewis (1991: 87) famously shied away from CAI just a few pages after introducing it. But, no matter where Lewis eventually ended up, it is clear that what initially attracted him to CAI was its apparent ability to make mereology ontologically innocent. Similar claims have been, and still are, made throughout the literature on CAI.²²

However, we will now argue that this is a mistake: CAI cannot deliver the ontological innocence of mereology without unduly restricting Plural Comprehension.

2.1 A limit of CAI

CAI gives a very clear sense to the claim that a fusion is nothing over and above its parts: a fusion is literally identical to its parts. However, the ontological innocence of mereology requires more than that. To see how CAI falls short, let’s return to our formulation of CAI from §1.1:

$$\text{CAI: } \forall x \forall y (Fxy \leftrightarrow (Ix \wedge x = y)).$$

This biconditional tells us when a given individual counts as a fusion of a given plurality: the individual is a fusion of the plurality iff the individual is identical to the plurality. However, it does not tell us whether any pluralities actually have any fusions. You cannot infer that any pluralities have fusions, because no guarantee has been given that any pluralities are identical to individuals. So, the existence of a fusion still appears to be something over and above the existence of its parts: a commitment to some parts is not yet a commitment to its fusion; to get from the first to the second, you need to add a commitment to an individual identical to the parts.

This observation is due to Cameron (2012).²³ But Cameron does not think that his observation undermines CAI’s ability to make mereology *ontologically* innocent:

Think of the view as follows: there is a fundamental property of *being an individual* that can be had by a collection of things, but need not be (at least, it’s *conceptually* possible that it not be had by some collection of things) [...] The *more that has to happen* for the Xs to compose is for the Xs to have this fundamental property. God has to do more to make the Xs compose than to make the Xs: He must make the Xs have this property. He must grant them individuality! If God gives the Xs this property He thereby ensures that there is an individual that has this property. But it simply doesn’t follow that He’s brought some new thing into being: the individual is identical to the collection that existed before and lacked the property. It’s just that these things are fundamentally a different way: they are now an individual whereas they weren’t before. (Cameron 2012: 550–1)

²² See for example: Sider 2007: 54–5; Cotnoir 2014: 7; Hawley 2014; Bennett 2015; Cotnoir and Varzi 2021: 195–6; Payton 2021b: 9194; Loss 2021a, 2021b: S4519–20, 2022. (But note that Loss and Cotnoir are primarily concerned with variations on CAI, which we discuss in §3.)

²³ McDaniel (2010) was the first to demonstrate that CAI was consistent with nihilism, but it was Cameron who made the point in the way presented above.

However, it is essential to remember that CAI-ists are working with the demanding conception of ontology, according to which only individuals really exist, in the sense that only $\exists x(Ix \wedge \phi x)$ expresses an *ontological* commitment; $\exists x(Px \wedge \phi x)$ expresses an *ideological* commitment about how individuals can be grouped together. So, given this conception of ontology, if God granted individuality to a plurality, They precisely *would* be bringing something new into being: before we had a merely ideological commitment to a plurality; now we have a genuinely ontological commitment to an individual.

It might be helpful to put the point like this. Lewis asked us to imagine drawing up an ontological inventory of what exists according to our scheme of things; he claimed that it would be double counting to include a fusion along with each of its parts, since the fusion is identical to its parts. But that is a mistake. It would be double counting to include the fusion along with the *plurality* of its parts. But a CAI-ist should not have included that plurality in advance of identifying it with a fusion: we are meant to be writing an *ontological* inventory of what *exists*, and according to CAI, a plurality exists only if it is identical to a fusion. What uncontroversially exists at the outset is each part, but relative to each of its parts, the fusion *is* something new: no exhaustive ontological inventory can afford to miss it out.²⁴

2.2 Building on CAI: reconceptualising commitments

If we are right, then CAI is not enough *by itself* to make mereology ontologically innocent. However, nobody ever said that the full CAI-ist picture of composition must follow straight from the mere statement that composition is identity. CAI-ists are free to build on CAI in their attempt to deliver the innocence of mereology.

What exactly would a CAI-ist need to add? Well, it is important to recognise that it would not be enough merely to add some principles saying that certain pluralities are also individuals. That would just be to make some ontological commitments! Rather, what they need to do is develop the philosophical picture behind the formalism of CAI in a way that makes the move from a plurality to its fusion seem innocent.

Here is our best attempt to develop such a picture. We are used to drawing a sharp distinction between a theory's ontology and its ideology. However, according to this CAI-ist picture, that is a mistake. Ontology and ideology are (at least sometimes) two sides of the same coin. When we commit ourselves to a plurality of individuals, we are making a merely ideological commitment about how individuals can be collected together: $\exists x(Px \wedge \phi)$. However, we can then choose to *reconceptualise* this commitment as an ontological commitment to the collection itself: $\exists y(Iy \wedge \phi)$. This reconceptualisation does not involve any change in *what* we are committed to. We are merely reconceiving of the *way* in which we are committed to it. In Lewis's (1991: 81) words, we are committing ourselves to the plurality 'all together' instead of 'one at a time'. So the individual we are now committed to is identical to the plurality we started with: $\exists x \exists y (Px \wedge Iy \wedge x = y \wedge \phi)$. In other words, we are now committed to a fusion. Moreover, it is always our free choice

²⁴ See Payton (2022) for a related point.

to reconceptualise an ideological commitment to a plurality as an ontological commitment to an individual: we are not relying on reality to provide an individual to vouchsafe the reconceptualisation; rather, the possibility of this reconceptualisation suffices for the existence of this individual. (Not that a fusion springs into existence when we first reconceptualise a commitment to its parts; the fusion is there all along because of the standing possibility of this reconceptualisation.)²⁵

This is certainly an unusual way of thinking about ontology and ideology. However, it is often remarked that the CAI-ist picture is radical — in Sider’s (2014: 211) words, it is ‘strange and strangely compelling’ — and this is our best attempt to articulate what this radical picture should involve. Moreover, this sort of picture is not unprecedented. Linnebo and Rayo (2012: esp. §7), for example, took a similar view about the ontological hierarchy of sets and the (supposedly) ideological hierarchy of types.

We are happy to grant, just for the sake of argument, that this CAI-ist picture of ontology and ideology is coherent. (Or, better, that this initial sketch of the picture can be developed into something coherent.) We also grant that, on this picture, mereology is ontologically innocent: an ontological commitment to a fusion is nothing more than an ideological commitment to a plurality of parts that we have chosen to look at in a different way. However, we do want to highlight one important feature of this picture: it requires that there be no *mere manys*, i.e. pluralities that are not also individuals.

If there are any mere manys, then we are not always free to reconceptualise our ideological commitment to a plurality as an ontological commitment to an individual. Rather, there are some pluralities that are ontologically privileged — they are individuals — and others that aren’t. And in that case, whenever we make the move from a plurality to an individual identical to it, we are holding ourselves hostage to ontological fortune: we cannot make this move unless reality cooperates by supplying an individual identical to the plurality we started with. This is just another way of saying that the move involves adopting a substantive ontological commitment.

We take this to be an essential requirement on CAI’s making mereology ontologically innocent: *there must be no mere manys*. In fact, we suspect that mere manys would really need to be *impossible*, but we will not pursue this stronger requirement. Just denying that mere manys *actually* exist is deeply problematic, because it runs head first into Cantor’s Theorem.

2.3 Cantor’s Theorem

Relative to CAI, the claim that there are no mere manys is equivalent to *Universalism* (or *Unrestricted Composition*):

Universalism: $\forall y \exists x Fxy$.

²⁵ To our knowledge, no CAI-ist has explicitly articulated this picture at any length. However, we think that many of the things that philosophers have said on behalf of CAI are suggestive of this picture, for example: Lewis 1991: 81–2; Sider 2007: 61; Bøhn 2014, 2016: 420.

Informally, Universalism states that, for any things, there is a fusion of those things. CAI has had its ups and downs with Universalism. In the early literature, it was widely thought that CAI implied Universalism.²⁶ However, since McDaniel (2010) and Cameron (2012), it is now generally accepted that this was a mistake: as we explained in §2.1, CAI does not by itself imply that any plurality has a fusion. Now, discovering that CAI does not imply Universalism may have been disappointing, since it would have been nice to have such a direct answer to the Special Composition Question, but that disappointment does not matter for our purposes. We are not dealing with CAI as an empty formalism, but as underwritten by a radical picture of ontology and ideology. And we have just argued that this radical picture does involve a commitment to Universalism.

Unfortunately, however, there is a bigger problem waiting in the wings. Far from *implying* Universalism, CAI threatens to be *inconsistent* with it. The problem is that, if there are at least three individuals, the claim that there are no mere manys violates Cantor's Theorem. Roughly, Cantor's Theorem tells us that, if there are at least three individuals, then there are more pluralities of individuals than individuals.²⁷ But, if there are no mere manys, then every plurality is an individual, and so there must be at least as many individuals as pluralities. Contradiction.²⁸

This is an informal statement of Cantor's Theorem. But it will be helpful later (§3.2) to be a bit more careful now. Formally, then, this is the version of Cantor's Theorem that concerns us here:

Cantor's Theorem: If there are at least three individuals, then there cannot be a dyadic relation, R , which is both:

- (i) *total* from pluralities to individuals, i.e. $\forall x(Px \rightarrow \exists y(Iy \wedge Rxy))$;
- (ii) *injective* from pluralities to individuals, i.e. $\forall x\forall y\forall z((Px \wedge Py \wedge Iz \wedge Rxz \wedge Ryz) \rightarrow x = y)$.

Clearly, identity is injective from pluralities to individuals. (Or, at least, that is clear if it makes sense to express identities between pluralities and individuals at all.) So, given that there at least three individuals, Cantor's Theorem implies that identity cannot be total from pluralities to individuals. Or in other words: there must be some mere manys.

Here, then, is a summary of our argument so far. CAI cannot make mereology ontologically innocent if there are any mere manys. But, assuming that there are at least three individuals, Cantor's Theorem implies that there are mere manys. (In

²⁶ See for example: Harte 2002: 114; Merricks 2005: 629–30. See also Sider 2007: 61–2.

²⁷ Cantor's Theorem fans may be surprised by the requirement that there be at least *three* individuals. Usually, the plural form of Cantor's Theorem applies as soon as there are at least two individuals. However, as we will shortly explain (§2.4), we are working with a version of Plural Comprehension that only allows us to introduce *proper pluralities*, i.e. pluralities of more than one individual. In this setting, Cantor's Theorem only comes into effect when there are at least three individuals.

²⁸ Note that this argument is structurally different from arguments due to Calosi (2016), Loss (2018), and Yi (2021) that CAI entails nihilism and is thus inconsistent with Universalism. Crucially, by basing our wider argument on Cantor's Theorem, we have made it far harder for CAI-ists to resist, as we explain in §3.

fact, if there are at least four individuals, then the *majority* of pluralities are mere manys!²⁹) So CAI cannot make mereology ontologically innocent.

2.4 Restricting Plural Comprehension

The most obvious way of trying to resist our argument is simply by rejecting Cantor's Theorem. However, Cantor's Theorem is a consequence of the following *Plural Comprehension* scheme (where $\phi_{y/x}$ is the result of substituting y for every free occurrence of x in ϕ):

$$\begin{aligned} \text{Plural Comprehension: } \exists x \exists y (x \neq y \wedge Ix \wedge Iy \wedge \phi \wedge \phi_{y/x}) \rightarrow \\ \exists y \forall x (x \varepsilon y \leftrightarrow (Ix \wedge \phi)), \\ \text{whenever } y \text{ does not appear in } \phi. \end{aligned}$$

Assume that there are at least three individuals. Now suppose for *reductio* that there is some relation R which is total and injective from pluralities to individuals. Let R^*x abbreviate $\exists z (Pz \wedge Rzx \wedge \neg x \varepsilon z)$, and consider this instance of Plural Comprehension:

$$\exists x \exists y (x \neq y \wedge Ix \wedge Iy \wedge R^*x \wedge R^*y) \rightarrow \exists y \forall x (x \varepsilon y \leftrightarrow (Ix \wedge R^*x)).$$

It is easy to show that the antecedent is implied by our assumption that there are at least three individuals,³¹ and so we can detach the consequent. Call the plurality introduced by this consequent a . Since R is total, there is some individual b such that Rab . Now ask whether $b \varepsilon a$. If $b \varepsilon a$, then R^*b ; however, since R is injective, a is the only plurality such that Rab , and so it follows that $\neg b \varepsilon a$. But if $\neg b \varepsilon a$, then R^*b , and so $b \varepsilon a$. Contradiction.

So, if a CAI-ist wanted to dodge Cantor's Theorem, they would have to either deny that there are at least three individuals, or restrict Plural Comprehension. We assume that no CAI-ist would want to deny that there are at least three individuals,³² and so their only option would be to restrict Plural Comprehension. But Plural Comprehension is a core part of standard plural logics. We rely on Plural Comprehension whenever we introduce *the* ϕ s, just on the grounds that at least two

²⁹ If there are κ individuals, then there are $2^\kappa - 1 - \kappa$ proper pluralities (i.e. excluding an empty plurality and any singleton pluralities). $2^\kappa - 1 - 2\kappa > \kappa$ whenever $\kappa \geq 4$.

³⁰ The entities introduced by this version of Plural Comprehension are not explicitly given as pluralities. However, that they are pluralities follows from the definition of P given in §1.1.

³¹ Proof: Continue to assume that there are at least three individuals — a, b, c — and that R is total and injective from pluralities to individuals. Plural Comprehension yields at least four pluralities: $[a, b, c], [a, b], [a, c], [b, c]$. Suppose that no individual satisfies R^* . In that case, R must injectively relate these four pluralities to just three individuals — a, b, c — which is impossible. Now suppose that exactly one individual satisfies R^* , and let it be a . In that case, R cannot relate $[a, b, c], [a, b]$ or $[a, c]$ to a ; and so R must injectively relate these three pluralities to just two individuals — b, c — which is also impossible.

³² More precisely, CAI would be trivial if nihilism were true, but if we rule out nihilism, CAI implies that there are at least three individuals. If nihilism is false, then there must be at least one fusion with two proper parts. (We are here assuming Weak Supplementation; see the end of §4.2.) CAI then implies that the fusion and its parts are all individuals.

individuals satisfy ϕ .³³ Any restriction on the scheme amounts to a reduction in the pluralities made available to us.

Still, CAI-ists would not be the first philosophers to restrict Plural Comprehension because of trouble with Cantor. However, we think that they are especially badly placed to *justify* their restriction, as we will now explain.

Intuitively, a plurality is just a definite collection of individuals. So, as Yablo (2006: 151–2) points out, there seem only to be two intelligible ways for an instance of Plural Comprehension to fail to define a plurality: it could be that ϕ is indefinite, in the sense that some individual neither definitely satisfies nor definitely fails to satisfy ϕ ; or it could be that we are testing ϕ against an indefinite domain of individuals. However, we will focus only on the second possibility, since we can safely assume that all other sources of indefiniteness (e.g. vagueness) have already been eliminated.

The idea that the domain of individuals might be indefinite is familiar from the philosophy of set theory. Start with any definite domain of individuals you like. Next, for each plurality of individuals in that domain, add the corresponding set.³⁴ By Cantor’s Theorem, this process must yield a proper extension of the domain that you started with. And, since this process can be applied to *any* definite domain, it follows that the domain of *all* individuals must be *indefinitely extensible*.³⁵

We do not want (or need) to take a stand here on whether set theory is best understood on the model of indefinite extensibility. What matters for our purposes is that even advocates of indefinite extensibility agree that every instance of Plural Comprehension should hold when the quantifiers range over a definite domain of individuals. (That is why they can still use Cantor’s Theorem.) After all, if we are testing a *definite* domain of individuals against a *definite* formula ϕ , then we should define a definite plurality. However, CAI-ists cannot retain Universalism without restricting Plural Comprehension even relative to certain definite domains of individuals.

Start with a definite domain of at least three individuals, D . Next, add in all of the fusions of individuals in D , and call the result D^+ . D^+ is a definite extension of D , and so is a definite domain. (Here is another way to see this: a domain is definite iff it has a cardinality;³⁶ and if κ is the cardinality of D , then D^+ has a cardinality no greater than $2^\kappa - 1$.) Universalism implies that every plurality of individuals in

³³ Payton (2021a) offers a novel analysis of ‘the ϕ s’: very roughly, ‘the ϕ s’ refers to every individual that has no part that does not overlap a ϕ . So, for example, ‘the bricks in the wall’ refers to a plurality that not only includes the bricks in the wall, but also the electrons that are parts of the bricks. However, *contra* Payton, this analysis does not correctly predict the behaviour of plural definite descriptions. To make this as vivid as possible, consider this description: ‘the things other than electrons in the wall’. Intuitively, this description should not refer to electrons, since they are explicitly ruled out. But, on Payton’s account, it still refers to them, because electrons are parts of other things in the wall.

³⁴ Plurality a corresponds to set b iff $\forall x(x \varepsilon a \leftrightarrow x \in b)$.

³⁵ See Fine 2006; Yablo 2006; Linnebo 2010, 2018b; Studd 2019; Florio and Linnebo 2021: ch. 11; Berry 2022.

³⁶ This follows from the common assumptions that a domain is definite iff it is set-sized, and that a domain is set-sized iff it has a cardinality. For related discussion, see Shapiro and Wright (2006). It is worth noting, though, that none of this amounts to an analysis of ‘definite’, since that notion was supposed to help *explain* when a domain forms a set (see Shapiro and Wright 2006: 265; Linnebo 2018a: 201).

D^+ has a fusion. But we can now be more specific. Every plurality of individuals in D^+ must have a fusion *in* D^+ : the fusion of some fusions is the fusion of the parts of those fusions, and so D^+ is already closed under fusion. It then follows from CAI that every plurality of individuals in D^+ is identical to an individual in D^+ . That would obviously contradict Cantor’s Theorem, and so CAI-ists cannot retain Universalism unless they restrict Plural Comprehension over D^+ .

This is what makes restricting Plural Comprehension particularly problematic for CAI-ists. They would have to claim that some instances fail to define a plurality *even when* ϕy is definite and the domain of individuals is definite. But how can that be, when a plurality *just is* a definite collection of individuals?

We are not aware of a satisfying answer to this question, and so we do not think that CAI-ists should restrict Plural Comprehension.³⁷ At the very least, we think that preserving unrestricted Plural Comprehension over definite domains is a desideratum that any attempt to make mereology ontologically innocent should ideally meet. The CAI-ist attempt does not satisfy this desideratum, but our own manyist attempt will (see §4).

3 Variations on CAI

Over the course of the last section, we argued that CAI, as formulated in §1.1, fails to make mereology ontologically innocent. However, a number of self-identifying CAI-ists have offered alternative formulations. In this section, we will look at two leading alternatives, and argue that neither of them can adequately deliver the innocence of mereology (§§3.1–3.2). We will then show that our argument can be generalised to cover any version of CAI whatsoever (§3.3).

³⁷ Bøhn (2016) argues that CAI blocks Cantor’s Theorem, and he acknowledges in a footnote (p. 420 fn. 21) that this means CAI-ists will have to restrict Plural Comprehension. He attempts to justify this restriction by claiming that pluralities do not include individuals *absolutely*, but only relative to concepts. However, that runs counter to the intuitive conception of a plurality as just a definite collection of individuals; and as we have seen, on that intuitive conception, it seems impossible to justify restricting Plural Comprehension over a definite domain. What is more, Bøhn does not offer any concrete details about how exactly Plural Comprehension should be restricted.

Payton (2021a: §6) also rejects Plural Comprehension, but he does provide a detailed alternative. Very roughly, Payton’s comprehension scheme states: if at least one thing is ϕ , then there is a plurality that includes x iff every part of x overlaps a ϕ . (Sider (2014: §3) and Saucedo (2025: §3) discuss related schemes.) However, Payton still does not explain how standard Plural Comprehension could intelligibly fail over a definite domain. (Saucedo (2025: §5) offers a metaphysical defence for his rejection of Cantor’s Theorem, but he also does not explain how a definite condition could fail to define a definite plurality over a definite domain.) Moreover, adopting Payton’s (or Sider’s or Saucedo’s) new comprehension scheme would lead to complications elsewhere. For example, pluralities are often used as the extensions of predicates, and, standardly, the extension of (e.g.) ‘is a brick’ should only include bricks; however, Payton’s comprehension scheme provides no such extension, since it never delivers a plurality that includes bricks without also including all the parts of those bricks. In recent work, Payton (2025) has developed a novel non-standard semantics that does not require the extension of ‘is a brick’ only to include bricks; but, as Payton (p. 1010) is the first to acknowledge, his semantic theory is fairly complex, which should count as a cost of his approach (and which our manyism will not incur).

3.1 Atomic Composition as Identity

Aficionados of the CAI literature might worry that some of §2 sounds like old news. We argued that CAI cannot deliver the innocence of mereology without restricting Plural Comprehension. But, some time ago, Sider (2007: §3, 2014: §2) demonstrated that (given some uncontroversial background assumptions) CAI implies the following Collapse principle:³⁸

$$\text{Collapse: } \forall x(Px \rightarrow \forall y(\exists z(Fzx \wedge y < z) \leftrightarrow y \varepsilon x)),$$

where $y < z$ symbolises *y is a proper part of z*. In plainer words, Collapse tells us that something is a proper part of the fusion of a given plurality iff it is included in that plurality.³⁹ This Collapse result has a number of bad consequences, but the worst is that it forces (non-trivial) CAI-ists to restrict Plural Comprehension (Sider 2014: §3).⁴⁰

Importantly, however, we are not merely rehashing Sider's old result. To show this, we will present Loss's (2021a,b, 2022) *Atomic Composition as Identity* (ACAI), and argue that, although ACAI avoids Collapse without restricting Plural Comprehension, it is still vulnerable to our argument.

According to the original version of CAI, a plurality's fusion is an individual identical to that plurality. Loss's ACAI restricts this thesis to pluralities of *mereological atoms*:

$$\text{ACAI: } \forall y(Ay \rightarrow \forall x(Fxy \leftrightarrow (Ix \wedge x = y))),$$

where Ay iff y is an individual atom or a plurality of atoms.⁴¹

What about pluralities that include non-atoms? Loss (2021b: §5) originally took a radical stance: he denied that there were any such pluralities. (So, for example, he denied that any plurality includes your favourite mug, since mugs aren't atoms.) However, this required imposing an untenable restriction on Plural Comprehension,

³⁸ For related arguments, see also: Yi 1999; Calosi 2016; Loss 2018.

³⁹ Here is how Sider put his argument in 2007. Let a be a plurality of individuals, and b be the fusion of that plurality. It should be trivial that everything included in a is a proper part of b , so all we need to show is that every proper part of b is included in a . Let c be a proper part of b . Whenever we fuse a fusion with one of its own parts, we just get the original fusion back, and so b should also be the fusion of $[b, c]$, i.e. the plurality of b and c . By CAI, it follows that $a = [b, c]$. So, since $c \varepsilon [b, c]$, it follows that $c \varepsilon a$.

As formulated, this argument tacitly relies on an instance of Plural Comprehension to deliver the plurality $[b, c]$. However, it would take an especially radical restriction on Plural Comprehension to prevent the introduction of a plurality of two given individuals. Moreover, Sider (2014: 212–3) has presented an alternative proof of Collapse that swaps Plural Comprehension for *Plural Covering*: $\forall x \forall y(x < y \rightarrow \exists z(Fyz \wedge x \varepsilon z))$.

⁴⁰ More precisely, Collapse forces any CAI-ist who rejects nihilism — which they should, since nihilism trivialises CAI — to restrict Plural Comprehension. If nihilism is false, then (assuming Weak Supplementation) there is at least one fusion, a , with two proper parts, b and c . By CAI, the fusion and its parts are all individuals. Plural Comprehension then implies that there is a plurality, $[a, b]$, which includes a and b but nothing else. However, c is a part of the fusion of $[a, b]$, since fusing a with one of its own parts just gives you the same initial fusion. Collapse therefore implies that $c \varepsilon [a, b]$. Contradiction.

⁴¹ This is not *quite* how Loss (2021a: 9203) formulates ACAI, but we take the differences to be incidental.

and in subsequent work, Loss (2021a, 2022) has stepped back from this extreme view. He now concedes that there are pluralities which include non-atoms, although he still maintains that pluralities of atoms are metaphysically privileged: according to Loss, these are the only pluralities that the most ‘joint-carving’ plural quantifier quantifies over. (Loss is here drawing on Sider’s (2011) idea that some quantifiers carve reality closer to its structural joints than others.) Crucially, however, Loss now accepts a fully unrestricted version of Plural Comprehension, albeit formulated with a less-than-maximally-joint-carving plural quantifier.

Now that Loss admits pluralities which include non-atoms, he needs to explain what their fusions are. But that is straightforward: the fusion of plurality a is an individual identical to the plurality of every atomic part (including improper parts) of every individual included in a . It follows that the fusion of a is identical to a iff a is an individual and also a plurality of atoms, as desired. (For related discussion, see: Loss 2021a: 9203.)

ACAI has two important virtues. First, it is immune to Sider’s Collapse argument.⁴² Second, ACAI is consistent with Universalism, even in the presence of unrestricted Plural Comprehension. Cantor’s Theorem still implies that there are mere manys, but relative to ACAI, that result no longer contradicts Universalism. Given ACAI, and assuming that there is no gunk,⁴³ all that is required for Universalism to be true is that every plurality of atoms be identical to an individual, and Cantor’s Theorem does not prohibit that.

However, despite being consistent with Universalism, ACAI does not dodge our argument from §2. We were never really concerned with the question of Universalism. Our concern was with the question of whether there are any mere manys. These two questions happen to be equivalent given CAI, but that was never the point. Now that we are working with ACAI, which does not retain this equivalence, we want to maintain our focus on mere manys.

Just like CAI, ACAI does not imply that any plurality has a fusion, since it supplies no guarantee that any plurality of atoms is identical to an individual. For this reason, ACAI is no more able than CAI to vindicate the innocence of mereology all by itself. We need to supplement it with a philosophical picture that makes the move from a plurality of atoms to an individual identical to that plurality seem innocent. Our best suggestion is the same picture that we offered in §2.2: the move is innocent because ideology and ontology are two sides of the same coin; we are free simply to choose to reconceptualise an ideological commitment to a plurality of atoms as an ontological commitment to an individual, without relying on any cooperation from reality.

However, that picture remains incompatible with the existence of mere manys. If there are mere manys, then that shows that we are not always free to reconceptualise an ideological commitment as an ontological commitment; that reconceptualisation

⁴² Sider’s argument relied on the fact that, if b is a fusion of plurality a , and c is a proper part of b , then CAI implies that $a = b = [b, c]$. But ACAI has no such implication: fusions are only ever identical to pluralities of atoms, and $[b, c]$ includes the non-atomic b . For exactly the same reason, ACAI is also immune to Calosi’s (2016), Loss’s (2018) and Yi’s (2021) arguments that CAI implies nihilism.

⁴³ In other words: assuming that everything is ultimately composed of mereological atoms.

is permissible only when reality happens to provide us with an appropriate individual, which is to say that it introduces a novel ontological commitment. Thus, Cantor's Theorem — which implies that there are mere manys — undermines ACAI's ability to deliver the ontological innocence of mereology.

This might not yet strike you as obviously correct. Given ACAI, wouldn't it be enough if no plurality of *atoms* were a mere many? However, it is important to remember that pluralities of atoms are not *conceptually* privileged. From a conceptual point of view, they are just some pluralities among many, no better and no worse. Loss's claim is only that pluralities of atoms are *metaphysically* privileged. But, of course, that is not something we can stipulate by fiat. Reality chooses which pluralities (if any) to privilege metaphysically, and we make our best conjectures about how it chose. So, if the move from a plurality to its fusion is permitted only when some appropriate plurality has been metaphysically privileged, then it still involves adopting a substantive ontological commitment: if reality cooperates, and bestows the required privilege, then we may introduce an ontological commitment to a fusion; otherwise, we must settle for an ideological commitment to a mere many.

To make this vivid, imagine that the only metaphysically privileged plurality is the universal plurality.⁴⁴ In that case, if a plurality is identical to an individual iff it is metaphysically privileged, then the only fusion would be the universal fusion. Now, we are not suggesting that an ACAI-ist needs to say anything to rule this possibility out. They have their theory about which pluralities are privileged, and that is good enough. However, the point remains that they make substantive claims about which pluralities are privileged, and hence which individuals exist. That is enough to show that, for an ACAI-ist, a fusion is an extra ontological commitment over and above its parts.

3.2 Composition as General Identity

CAI was formulated in terms of numerical identity — $x = y$. However, according to Cotnoir (2013), numerical identity is just a special case of *general identity* — $x \approx y$ — and we should trade-in CAI for *Composition as General Identity* (CAGI):

$$\text{CAGI: } \forall x \forall y (Fxy \leftrightarrow (Ix \wedge x \approx y)).$$

General identity is meant to capture the idea that x is the same 'portion of reality' as y , while allowing that x and y might carve up that portion in different ways: x might carve it up as five trees, while y carves it up as one copse. Cotnoir (2013: §2) provides a set-theoretic model of the behaviour of \approx , but we do not need to break open that formal machinery here. Instead, it will suffice for our purposes to rest on our intuitive grasp of the idea that we can carve up one portion of reality in different ways.

What matters now is that Cotnoir (2013: 304) takes numerical identity to be stronger than general identity:

⁴⁴ See Saucedo (2022: §4.1) for related discussion.

$$\forall x \forall y (x = y \leftrightarrow (x \sim y \wedge x \approx y)).$$

The new relation, $x \sim y$, is meant to express the idea that there is some *count* to which x and y both belong. Very roughly, a *count* of a given portion of reality is a way of carving that portion up into non-overlapping sub-portions. For Cotnoir, then, numerical identity is just *intra-count* general identity. To see that this is stronger than general identity *simpliciter*, let a be a plurality of atoms, b be a plurality of five trees composed of those atoms, and c be a copse composed of those trees. Given CAGI, we have $a \approx b \approx c$,⁴⁵ but a , b and c are pairwise numerically distinct, since they all belong to different counts (i.e. to different ways of carving up one and the same portion of reality).

Here is another way of making pretty much the same point: general identity is *not* injective from pluralities to individuals; in the case above, $a \approx c$ and $b \approx c$, even though $a \neq b$. Consequently, it does not violate Cantor's Theorem to say that every plurality is *generally identical* to an individual. Cantor's Theorem states that, if there are at least three individuals, then no relation can be both total and injective from pluralities to individuals. So, since general identity is not injective, there is no formal barrier to its being total. It seems, then, that a CAGI-ist has the resources to respond to our argument from §2 by denying that there are any mere manys in the sense that they care about (i.e. pluralities which are not generally identical to individuals), without restricting Plural Comprehension.⁴⁶ And this prospect is only made more appealing by the fact that CAGI also avoids Sider's Collapse result (Cotnoir 2013: §4.1).⁴⁷

However, CAGI only has a hope of making mereology ontologically innocent if general identity really is a form of identity, and we cannot see how a non-injective relation could possibly be a form of identity.⁴⁸ Cotnoir (2013: §3) is very sensitive to this kind of worry. He attempts to soothe it by arguing that general identity satisfies a completely general version of Leibniz's Law:

Leibniz's Law[~]: $\forall x \forall y (x \approx y \rightarrow (\phi \leftrightarrow \psi))$, whenever ψ is a result of substituting a free occurrence of y for a free occurrence of x in ϕ .

On the face of it, CAGI provides some fairly obvious counter-examples to Leibniz's Law[~]. Go back to our earlier example, where a is a plurality of atoms, and b is a plurality of five trees composed of those atoms. Given CAGI, $a \approx b$, but it seems

⁴⁵ We assume here and throughout that \approx is an equivalence relation.

⁴⁶ We should mention that Cotnoir (2013: 299 fn.10 & 303 fn.15) does restrict Plural Comprehension, and mentions Cantor's Theorem as one of his justifications for this restriction. But, as we have just explained, Cantor's Theorem is no justification for a CAGI-ist to restrict Plural Comprehension. At any rate, in this subsection we are considering how a CAGI-ist might respond to our argument, even if that CAGI-ist is not Cotnoir himself.

⁴⁷ Although Cotnoir does restrict Plural Comprehension (see fn.46), that is not what blocks Sider's Collapse argument. (Cotnoir (2013: 314) himself is clear on this.) Instead, Sider's argument is blocked precisely by the fact that general identity is not injective from pluralities to individuals. Sider's argument relies on the fact that $c \varepsilon a$ follows from $a = b = [b, c]$; but, Cotnoir argues, $c \varepsilon a$ does not follow from $a \approx b \approx [b, c]$. Just as for ACAI, it also follows that CAGI is immune to Calosi's (2016), Loss's (2018) and Yi's (2021) arguments that CAI implies nihilism.

⁴⁸ Carrara and Lando (2016) offer an alternative argument that general identity is not a form of identity, on the grounds that general identity is not appropriately related to co-referentiality.

clear that various claims are true of a and not b : a includes atoms, b does not; a includes more than five things, b does not; and so on. However, Cotnoir (2013: 311–2) deals with these cases by suggesting that some predications are count-relative. The pluralities a and b are the same portion of reality, counted in different ways. So, relative to one count, b includes trees, but we can *recount* b as a plurality of atoms; relative to *that* count, b includes atoms. Similarly, relative to some counts, b only includes five individuals, but relative to others it includes more. (And of course, exactly the same goes for a .) So, $a \approx b$ only appears to provide counter-examples to Leibniz’s Law⁴⁹ when we evaluate ϕa and ϕb relative to different counts.⁴⁹

We have a general worry about this strategy. We suspect that to say that ϕx is true relative to some count is really just to say that $\exists y(x \approx y \wedge \phi y)$. And if that is right, then Cotnoir can’t really be saving Leibniz’s Law⁵⁰ in any substantial sense.⁵⁰ However, we do not want to pursue this worry here. That’s because, even without working through the full details of Cotnoir’s (2013: 308–13) formal account, we can already be sure that he hasn’t managed to preserve Leibniz’s Law⁵⁰ in its full generality. To see why, consider the following instance (with a and b as above):

$$a \approx b \rightarrow (a = a \leftrightarrow a = b).$$

By assumption, $a \approx b$ but $a \neq b$, since a and b were meant to illustrate Cotnoir’s idea that numerical identity is more demanding than general identity. So, since $a \approx b$, we can detach the consequent from the above conditional. But trivially $a = a$, and so $a = b$. Contradiction.

It is important to recognise that Cotnoir cannot use his machinery of count-relativity to deal with this counter-example. To begin with, it is clear that Cotnoir did not intend numerical identity to be count-relative. Recall that he thinks of numerical identity as intra-count general identity: $a = b$ iff $a \sim b \wedge a \approx b$. Cotnoir (2013: 311) is explicit that $a \approx b$ is not count-relative, and $a \sim b$ should not be either: $a \sim b$ is the claim that a and b belong to the same count; this should be an absolute claim *about* the counts relative to which other predications are evaluated. Moreover, even setting Cotnoir’s intentions to one side, little sense can be made of the idea that numerical identity is count-relative. Since a is the same portion of reality as b , no single way of recounting that portion could ever distinguish between them. So, relative to *every* way of counting that portion, $a = b$ would be true.

Cotnoir’s CAGI must, then, involve some restriction on Leibniz’s Law⁵⁰. This restriction may not need to be as severe as it would seem at first glance, thanks to Cotnoir’s machinery of count-relativity. But some restriction is needed if general identity is to be weaker than numerical identity. Otherwise, general identity will imply numerical identity, and so general identity will be injective from pluralities to individuals after all.

Now, we have to admit that the status of Leibniz’s Law is not entirely settled. The consensus is that Leibniz’s Law is sacrosanct — if a relation doesn’t satisfy

⁴⁹ Cotnoir (2013: §3.2) actually proposes two strategies for preserving Leibniz’s Law⁵⁰: one is to introduce a count-relativity to our predications; the other is to introduce a subvaluational semantics. We have focused on the count-sensitivity proposal, but all of our points carry over straightforwardly to the subvaluational proposal.

⁵⁰ For related discussion, see Hawley (2013: §2) on ‘antipodean counterparts’.

Leibniz's Law, then it just isn't *identity* — but that consensus does not amount to total unanimity.⁵¹ It is probably best, then, to think of maintaining Leibniz's Law as a second desideratum on the attempt to make mereology ontologically innocent: we should not restrict Leibniz's Law for the sense of identity in which fusions are 'identical' to their parts; otherwise, we will open up real doubts about whether we are genuinely *identifying* fusions with their parts. Cotnoir's CAGI does not satisfy this desideratum, but as we will shortly see, our manyism does.

3.3 The argument generalised

So far, we have looked at three particular versions of CAI — original, ACAI, and CAGI — and argued that none of them can adequately deliver the ontological innocence of mereology. But our argument is really wholly general. The guiding idea behind CAI is that a commitment to a fusion is an ideological commitment to a plurality, reconceptualised as an ontological commitment to an individual. But no theory built on this idea can make mereology ontologically innocent without violating one of these desiderata:

Desideratum 1: Do not restrict Plural Comprehension, at least over definite domains.

Desideratum 2: Do not restrict Leibniz's Law for the sense of identity in which a fusion is 'identical' to its parts.

Let D^+ be a definite domain of at least three individuals that is already closed under fusion. (As we explained in §2.4, closing a definite domain under fusion yields a definite domain.) If we do not restrict Plural Comprehension or Leibniz's Law over D^+ , we can use Cantor's Theorem to show that some pluralities defined over D^+ are not identical to any individuals in D^+ . But it then follows that these pluralities are mere manys: if they were going to be identical to any individuals, they would have to be identical to fusions already in D^+ . This follows because, as Sider (2007: 60 fn. 24) points out, the right-to-left direction of CAI is a trivial consequence of Leibniz's Law and the fact that every individual is a (possibly improper) fusion of itself:

$$\forall x \forall y ((Ix \wedge x = y) \rightarrow Fxy).^{52}$$

So, if a plurality of individuals in D^+ were itself an individual, then it would be a fusion of the original individuals, and would therefore already be included in D^+ . Thus, some pluralities defined over D^+ must be mere manys. And if there are any mere manys, then reconceptualising an ideological commitment to a plurality as an ontological commitment to an individual cannot be innocent, since it sometimes fails.

⁵¹ For versions of CAI that reject Leibniz's Law for the sense of 'identity' in which fusions are identified with their parts, see Baxter (1988, 2014), Turner (2014), and Bricker (2016). Bricker (2021) also considers an approach that is in many ways similar to Cotnoir's, but which involves unapologetically restricting Leibniz's Law; however, Bricker ultimately concludes that this approach is inadequate.

⁵² Suppose that x is an individual identical to y ; since Fxx , Leibniz's Law implies Fxy . (We have formulated the right-to-left direction of CAI as a claim about numerical identity; but, of course, exactly the same reasoning will apply for any other kind of identity, so long as it satisfies Leibniz's Law.)

4 How manyism delivers ontological innocence

In this section, we will argue that manyism can deliver the ontological innocence of mereology without violating either of the desiderata identified in §3.3. We will start (§4.1) by filling out the manyist picture in a little more detail. Then (§4.2) we will present a formalisation of manyism, and state our crucial *mereologist* result: manyism transforms all of the axioms of Classical Mereology into theorems of Super-Plural Logic. (Proofs are supplied in the Appendix.) Finally (§4.3) we use this formal result to argue that manyism makes mereology ontologically innocent.

4.1 Manyism: some missing details

CAI-ists think that a plurality is a fusion iff it is not a *mere many*. Manyists disagree. They think that *every* plurality of individuals is both a fusion and a *mere many*. Indeed, they think that a fusion *just is* a plurality of individuals, no more or less.⁵³ In this subsection, we will develop the manyist picture in more detail.

We would like to begin by admitting the obvious. Manyism is metaphysically revisionary: philosophers usually assume that fusions are individuals. It might also be semantically revisionary, insofar as we usually take names for fusions to be semantically singular, in the sense of referring to an individual. However, these revisions are far less drastic than they might seem to be, when we remember that manyists are operating with the liberal conception of ontology, according to which individuals and pluralities all exist in the very same sense. So manyists do not deny that fusions exist; they only deny that fusions are individuals. They also do not deny that a name for a fusion refers to exactly one existent; it's just that that existent is a plurality, not an individual. (So, for example, there is a perfectly ordinary sense in which there might be *exactly one* mug in your domain, despite the fact that mugs are mere many, namely: $\exists x \forall y (\text{Mug}(y) \leftrightarrow x = y)$.)

Manyism has another controversial consequence. If fusions are pluralities of individuals, and no individual is a plurality, then individuals must be mereological atoms, and every fusion must ultimately be composed of atoms.⁵⁴ So manyism implies that there is no 'mereological gunk'. We grant that this is a cost of manyism. However, there are a number of well-known strategies for minimising this cost.⁵⁵ Moreover, there is always the option of *ponens*-ing instead of *tollens*-ing. Manyism offers a clear conception of mereological fusion that forbids gunk: fusions are just pluralities of individuals, and individuals are not themselves pluralities. Rather

⁵³ You could, in principle, imagine a *restrictivist* manyist, who claimed that *some*, but not *all*, pluralities of individuals count as fusions. However, we find it hard to see what would motivate this restrictivist brand of manyism, since it is unclear why one *mere many* would be privileged over another. Moreover, our aim is to show that manyism *can* be developed in a way that secures the ontological innocence of mereology, and so we have chosen to focus on the version of manyism that stands the best chance. As we will shortly explain, our manyism takes every plurality of individuals to be a fusion as a matter of definition.

⁵⁴ It might be possible to revise manyism in a way that avoids this result, if we were willing to adopt a non-well-founded plural logic. (For related discussion, see Werner 2022.) However, we will set that possibility aside. Part of what we find so appealing about manyism is that, unlike CAI, it requires no logical innovations.

⁵⁵ See Williams 2006; Cameron 2007: 101–2; Sider 2013: 270–82.

than seeing this as a shortcoming of manyism, we could choose to see it as an argument for rejecting gunk.⁵⁶

According to manyism, then, a mereological atom is any individual, and a fusion is any plurality of individuals. Now, as well as pluralities of individuals, we also want to be able to talk about pluralities of fusions. That will be essential for formulating various mereological principles, such as Universalism: *for any plurality of fusions and atoms*, there is a fusion of that plurality. These pluralities of fusions will be *super-pluralities*, i.e. pluralities that are to pluralities of individuals as pluralities of individuals are to individuals.

Super-pluralities are sometimes looked on with a little bit of suspicion. Most of this suspicion stems from the fact that natural language doesn't include any uncontroversial examples of super-plural terms. However, there are still some decent candidates. For example, it is natural to identify bands with their members. (So, for example, it is natural to think of 'The Beatles' as a plural term that refers to John, Paul, George and Ringo.) But, if that is right, then plural quantification over bands (e.g. 'The bands she represents have made her a lot of money') is really a form of super-plural quantification.⁵⁷

But even if you remain unconvinced that natural languages include any super-plural quantification, all that really matters is that this quantification be *intelligible*. It is hard to know how to convince someone that a notion is intelligible, other than by showing them how it can be used. We do just that in our Appendix, where we derive mereologism from manyism. However, in the meantime, a mere appeal to authority might tide us over: many of the leading plural logicians are advocates of super-plural quantification.⁵⁸

So, a mereological atom is an individual, a fusion is a plurality of individuals, and a plurality of fusions is a super-plurality. Importantly, if no plurality is an individual — and the whole point of manyism is to avoid identifying pluralities with individuals — then no super-plurality is a plurality of individuals. Consequently, according to manyism, no super-plurality is a fusion. It follows that a fusion is identical only to its *atomic* parts. Take, for example, your favourite mug. According to manyism, that mug is identical to its atomic parts, but not to its handle and bowl. The handle and bowl are still parts of the mug, the mug is still their fusion, but the mug is not identical to them. That's because *they* are a super-plurality, and the mug is a plurality of individuals.⁵⁹

⁵⁶ For related discussion, see Thunder 2023: §2.3.3.2.

⁵⁷ See Hewitt (2012: 866 fn. 24) for a brief defence of the treatment of band names as plural terms. See Grimaud (2021: §5.1) for further plausible examples of super-plural terms in English, Icelandic, Finnish, and Khamtanga; see also Linnebo and Nicolas (2008) and Florio and Linnebo (2021: ch. 9).

⁵⁸ For example, see: Rayo 2006; Oliver and Smiley 2016: §8.4 & ch. 15; Florio and Linnebo 2021: esp. ch. 9.

⁵⁹ We can, then, think of manyism as what you get if you take Loss's ACAI, and remove the idea that fusions are individuals. Importantly, this is a revision to Thunder's (2023: 21) original version of manyism. It also distinguishes manyism as developed here from Carrara and Lando's (2016: 136–8) sympathetic reworking of Cotnoir's CAGI in super-plural terms: according to their reworked CAGI, a mug is (in some sense) identical to the plurality of its bowl and handle, even though the bowl and handle are themselves pluralities. Carrara and Lando motivate their revision of CAGI by claiming that, if a plural term, *a*, refers to some individuals, and a super-plural term, *b*, refers to pluralities whose union is exactly the original individuals, then *a* and *b* co-refer. We should emphasise that this

4.2 Formalising manyism

We are now in a position to begin formalising manyism. The background logic is (what we will call) *Super-Plural Logic* (SPL). This logic is designed to be a fragment of Oliver and Smiley's (2016: ch. 15) *Higher-Level Plural Logic*, with quantifiers restricted to individuals, pluralities of individuals, and pluralities of pluralities of individuals (i.e. super-pluralities).

SPL is single-sorted, and any term may take an individual, a plurality, or a super-plurality as value. The formation rules for SPL are exactly the same as the formation rules for standard first-order logic. We also port over all of the standard axioms and inference rules from first-order logic, including Leibniz's Law:

Leibniz's Law: $\forall x \forall y (x = y \rightarrow (\phi \leftrightarrow \psi))$, whenever ψ is a result of substituting a free occurrence of y for a free occurrence of x in ϕ .

However, SPL adds one extra basic logical predicate, ε , which expresses vertical inclusion. We then have various defined notions:⁶⁰

$$\begin{aligned} Ia &:= \forall x \neg x \varepsilon a \\ Ma &:= \forall x (x \varepsilon a \rightarrow Ix) \\ a \varepsilon^* b &:= a \varepsilon b \vee (Ib \wedge a = b) \end{aligned}$$

I is the manyist definition of individuals; it suffices for manyists because they don't believe that individuals can be pluralities (and they also don't believe in an empty plurality). M is for *mereological*, since, according to manyism, only individuals and pluralities of individuals can be related by the part-whole relation. The final defined notion, ε^* , is a useful device for avoiding singleton pluralities: if a is an individual, then $x \varepsilon^* a$ iff $x = a$, and so a can stand in for its own singleton plurality.

If we wanted to accommodate higher levels of super-pluralities, we would need a rich set of axioms governing ε . But we only need to go as high as pluralities of pluralities of individuals, and so these will do as our SPL-axioms for ε :

$$\begin{aligned} \text{Extensionality: } & \forall x \forall y ((\neg Ix \vee \neg Iy) \rightarrow (x = y \leftrightarrow \forall z (z \varepsilon x \leftrightarrow z \varepsilon y))) \\ \text{Plurality: } & \forall x \forall y (x \varepsilon y \rightarrow \exists z (z \neq x \wedge z \varepsilon y)) \\ \text{Super-Plurality: } & \forall x \forall y (x \varepsilon y \rightarrow Mx) \\ \text{Comprehension: } & \exists x \exists y (x \neq y \wedge Mx \wedge My \wedge \phi \wedge \phi_{y/x}) \rightarrow \\ & \quad \exists y \forall x (x \varepsilon y \leftrightarrow (Mx \wedge \phi)), \\ & \quad \text{whenever } y \text{ does not appear in } \phi \end{aligned}$$

Extensionality is a standard axiom of plural logic. Plurality ensures that pluralities always include more than one thing. (So there are no 'singleton' pluralities. Every plurality is a *many*.) Super-Plurality makes sure we don't go above pluralities of

is *not* how we are thinking about pluralities and super-pluralities. In our setting, no term referring to a plurality of individuals co-refers with a super-plural term, by a straightforward application of Leibniz's Law.

⁶⁰ When applying a definition, bound variables should be re-lettered as necessary to ensure that no variable free in the definiendum is bound in the definiens. For example, the definiens of Ix is (an alphabetic variant of) $\forall y \neg y \varepsilon x$, not $\forall x \neg x \varepsilon x$.

pluralities of individuals. (It also ensures that ε is well-founded.) In the presence of Super-Plurality, Comprehension gives us all the comprehension that we need. This new Comprehension scheme implies Plural Comprehension from §2.4, since Ix implies Mx . So Comprehension also implies the plural form of Cantor's Theorem from §2.3: there are more pluralities of individuals than individuals. (There are also more super-pluralities than pluralities of individuals.) But this theorem no longer poses any problems, since no individual is a plurality (and no plurality of individuals is a super-plurality, either).⁶¹

With this background in place, we can now formalise manyism. Crucially, we can think of manyism as providing us with definitions of mereological notions in super-plural terms. These definitions are not intended to perfectly capture our existing mereological notions, just as they are. (We are not claiming that manyism itself is analytic.) They are instead offered as *explications*, i.e. as precisifications of those notions.⁶² We are not *forced* to accept these manyist explications; but, as we will argue in §4.3, accepting them does yield some important philosophical benefits.

We start with a definition of *fusion* (including improper cases). Here is what a manyist intuitively wants: if a is an individual, then the fusion of a should just be a itself; otherwise, the fusion of a should be the plurality of individuals that a is ultimately constructed from. (For example, if a, \dots, e are individuals, then the fusion of the super-plurality $[[a, b, c], [c, d], e]$ should be the first-level plurality $[a, b, c, d, e]$.) And here is a particularly neat way of putting that into a definition:

$$Fab := \forall x(x \varepsilon^* a \leftrightarrow \exists y(x \varepsilon^* y \wedge y \varepsilon^* b)).$$

(We demonstrate that this definition is adequate in the Appendix.)⁶³ Next we need a definition of *parthood* (also including improper cases). There are several well-known ways of defining parthood in terms of fusion, and we have chosen this one:

$$a \leq b := \exists x(Fbx \wedge a \varepsilon^* x).⁶⁴$$

Of course, we still have to demonstrate that this definition of parthood is appropriately related to our earlier definition of fusion. The crucial result is as follows:

$$Ma \rightarrow (Fab \leftrightarrow (\forall x(x \varepsilon^* b \rightarrow x \leq a) \wedge \forall x(x \leq a \rightarrow \exists y(y \varepsilon^* b \wedge x \circ y))))),$$

⁶¹ If you are worried about indefinite extensibility, feel free to read the quantifiers in Comprehension as ranging over definite totalities. See footnotes 66 and 67 for further discussion.

⁶² See Carnap 1945: 513, 1950: ch. 1.

⁶³ It is also worth noting that this treatment of fusion can also make good sense of claims which intuitively describe the composition of multiple fusions, such as 'The molecules compose the cells'. This sentence is true iff both (i) each of the cells is a fusion of some of the molecules, and (ii) each of the molecules is fused into at least one of the cells. Letting c be the cells and m be the molecules, our account therefore yields the following: 'The molecules compose the cells' is true iff (i) $\forall x(x \varepsilon^* c \rightarrow \exists y \forall z((z \varepsilon^* y \rightarrow z \varepsilon^* m) \wedge Fxy))$, and (ii) $\forall x(x \varepsilon^* m \rightarrow \exists y(x \varepsilon^* y \wedge \exists z(Fzy \wedge z \varepsilon^* c)))$. This account can be extended straightforwardly to cover more complex claims, such as 'Some atoms compose some molecules, which compose some cells, which compose ... which compose an organism'.

⁶⁴ This is a *collective* notion of parthood. We can also define a *distributive* notion: $\forall z(z \varepsilon^* a \rightarrow z \leq b)$. If Ma , then a is collectively part of b iff a is distributively part of b . However, if $\neg Ma$, then distributive and collective parthood come apart: a is never collectively part of anything, but is distributively part of everything that its fusion is part of.

where $x \circ y$ abbreviates $\exists z(z \leq x \wedge z \leq y)$. We provide the necessary proof in the Appendix. We then go on to prove that, given the manyist definitions of fusion and parthood, all of the axioms of Classical Mereology (CM) become theorems of SPL:⁶⁵

Transitivity: $\forall x \forall y \forall z ((x \leq y \wedge y \leq z) \rightarrow x \leq z)$

Universalism: $\forall x \exists y Fyx$

Weak Supplementation: $\forall x \forall y (x < y \rightarrow \exists z (z \leq y \wedge \neg x \circ z))$

4.3 Mereologism and the innocence of mereology

The results described in the previous subsection have some important consequences. First, they guarantee that manyism does not lead to any unusual behaviour in our mereology: assuming that the axioms of SPL are all true, manyism implies that CM is true.⁶⁶ Second, they also guarantee that manyism does not lead to any unusual behaviour in our plural logic: given manyism, we do not need to add any mereological axioms to our plural logic; we only need to add some mereological definitions. (In particular, then, manyism must be immune to Sider's Collapse argument.) Third, we claim that they show that manyism successfully makes mereology ontologically innocent, without restricting Comprehension or Leibniz's Law.

We make this third claim because we think that SPL is logic, and so the manyist proof that SPL implies CM is a proof that CM is also logic.⁶⁷ We call this result *mereologism*.⁶⁸ Manyism delivers the ontological innocence of mereology precisely by delivering mereologism: a fusion is nothing over and above its parts in the sense that the existence of the parts *logically entails* the existence of the fusion; a commitment to some parts *logically entails* a commitment to their fusion.

Of course, this argument only works if you think that SPL is logic. Otherwise, the manyist proof that CM follows from SPL would not be a proof of mereologism. And there is room for resistance here. To begin with, you might be resistant to the whole idea of a plural logic. Or, less radically, you might be resistant to the idea of a *single-sorted* plural logic. Less radically still, you might resist the introduction of super-pluralities. However, our main aim in this paper is to recommend manyism as an alternative to CAI, and none of the above is a line that a CAI-ist could draw

⁶⁵ For a demonstration that the following is an axiomatisation of CM, see Hodva 2009. Hossack (2000: §5) gestures towards a similar result to the one we give in the Appendix, but with an important difference: Hossack's version of *Universalism* only states that every *plurality of individuals* has a fusion, whereas ours also implies that every *super-plurality* has a fusion. (Moreover, Hossack adopts the demanding conception of ontology, and so takes himself to be *eliminating* mereological fusions.) For additional related results, see Florio and Linnebo 2021: ch. 5.

⁶⁶ Or, if you are restricting Comprehension to definite domains, then manyism implies that CM is true relative to any definite domain of individuals. We suspect that this is the best that a believer in indefinite extensibility can hope for.

⁶⁷ Or, if you are restricting Comprehension to definite domains, then CM is logic, and so ontologically innocent relative to any definite domain of individuals.

⁶⁸ Russell (2017) presents an alternative version of mereologism (or something near enough), based on neo-Fregean abstraction. Russell introduces abstraction in order to preserve the orthodoxy that fusions are individuals. However, as Russell himself emphasises, the innocence of any given abstraction principle is always a vexed issue. (Plus, see Trueman 2014 for general doubts about the whole abstractionist programme.) What we have shown is that, by giving up on orthodoxy, manyists gain access to a much more straightforward form of mereologism.

in the sand: CAI-ists also rely on a single-sorted plural logic; and they also permit super-pluralities, since they grant that fusions are pluralities that can be included in pluralities.⁶⁹ So, is there any objection that a CAI-ist could offer to the manyist proof of mereologicism?

We think that this is a CAI-ist's best bet. Logic is meant to be ontologically innocent. Indeed, we manyists are trying to exploit that innocence: we claim that mereology is ontologically innocent precisely because we claim to have proven that CM is logic. But why should we think of SPL as ontologically innocent? After all, the Comprehension scheme introduces a commitment to pluralities and super-pluralities. Now, a CAI-ist has an answer to this question. They accept the *demanding* conception of ontology, according to which only individuals exist. So, for a CAI-ist, Comprehension merely introduces some ideological commitments. But manyists accept the *liberal* conception of ontology, according to which individuals, pluralities and (we now add) super-pluralities all exist in the same sense. So, doesn't it follow that, for a manyist, Comprehension introduces novel *ontological* commitments, and thus that SPL is not a form of logic after all?

This is an important challenge, but we think we can meet it. The demanding conception of ontology draws a distinction between two kinds of commitment: genuinely ontological commitment to individuals, and merely ideological commitment to pluralities. On the face of it, there is more than one way of rejecting this distinction. We could insist that pluralities exist in just the way that the demanding conception says that individuals exist. But, alternatively, we could insist that a commitment to an individual is no more substantive than a commitment to a plurality. We prefer the second option. We do not think that there is an especially robust kind of commitment, *genuinely* ontological commitment, against which we can contrast *merely* ideological commitment. Rather, we think that all there is to an ontological commitment is the kind of commitment that the demanding conception dismisses as ideological. In other words, we have not got to our liberal conception by starting with the demanding conception and *upgrading* our commitment to pluralities, but by *downgrading* our commitment to individuals.

There is, then, no substantial disagreement between us and the CAI-ists over how to understand our commitment to pluralities. (We label it 'ontological', they label it 'ideological', but that is just terminology. We agree about what the commitment amounts to.) So if CAI-ists are happy to grant that Comprehension is ontologically innocent by their lights, then they must grant that it is innocent by our lights too. Our disagreement is about individuals, not pluralities: we deny that individuals exist in any more demanding a sense than do pluralities. You might think that this sounds a bit deflationary, but that is not how we hear it. We just really believe in Quine's criterion: to be is to be the value of a variable, no more and no less.

⁶⁹ A CAI-ist might counter that they still only believe in first-level pluralities of individuals, since their fusions are also individuals. However, this reply is only worth making if there is an important ontological difference between individuals and non-individuals, which brings us to the objection we are about to discuss.

Appendix

In this appendix, we will prove that, given the manyist definitions of fusion and parthood, the axioms of Classical Mereology all become theorems of Super-Plural Logic.

A Super-Plural Logic

First, a quick refresher on Super-Plural Logic (SPL). SPL is single-sorted, and the formation rules are exactly the same as the formation rules for standard first-order logic. We also port over all of the standard axioms and inference rules from first-order logic, including Leibniz's Law:

Leibniz's Law: $\forall x \forall y (x = y \rightarrow (\phi \leftrightarrow \psi))$, whenever ψ is a result of substituting a free occurrence of y for a free occurrence of x in ϕ .

However, we also add one extra basic logical constant, ε , which expresses vertical inclusion.

Definition 1: We have various defined notions:

- (a) $a \sqsubseteq b := \forall x (x \varepsilon a \rightarrow x \varepsilon b)$
- (b) $a \sqsubset b := a \sqsubseteq b \wedge a \neq b$
- (c) $Ia := \forall x \neg x \varepsilon a$
- (d) $Ma := \forall x (x \varepsilon a \rightarrow Ix)$
- (e) $Pa := Ma \wedge \neg Ia$
- (f) $Sa := \exists x \exists y (x \varepsilon y \wedge y \varepsilon a)$
- (g) $a \varepsilon^* b := a \varepsilon b \vee (Ib \wedge a = b)$

(When applying a definition, bound variables should be re-lettered as necessary to ensure that no variable free in the definiendum is bound in the definiens.) This list is slightly longer list than we gave in §4.2. \sqsubseteq is horizontal inclusion, and \sqsubset is proper horizontal inclusion. I , P and S are *individual*, *plurality of individuals* and *super-plurality*. M is for *mereological*. And, as before, ε^* is a useful abbreviation for avoiding singleton pluralities.

Here are the SPL-axioms governing ε (where $\phi_{y/x}$ is the result of substituting y for every free occurrence of x in ϕ):

Extensionality: $\forall x \forall y ((\neg Ix \vee \neg Iy) \rightarrow (x = y \leftrightarrow \forall z (z \varepsilon x \leftrightarrow z \varepsilon y)))$

Plurality: $\forall x \forall y (x \varepsilon y \rightarrow \exists z (z \neq x \wedge z \varepsilon y))$

Super-Plurality: $\forall x \forall y (x \varepsilon y \rightarrow Mx)$

Comprehension: $\exists x \exists y (x \neq y \wedge Mx \wedge My \wedge \phi \wedge \phi_{y/x}) \rightarrow$
 $\exists y \forall x (x \varepsilon y \leftrightarrow (Mx \wedge \phi)),$
 whenever y does not appear in ϕ

B Introducing fusion

All of the basic terms of mereology will be defined in super-plural terms. We start with *fusion* (including improper cases):

Definition 2: $Fab := \forall x(x \varepsilon^* a \leftrightarrow \exists y(x \varepsilon^* y \wedge y \varepsilon^* b))$

Lemma 3: A list of elementary facts about F :

- (a) $Mb \rightarrow (Fab \leftrightarrow a = b)$
- (b) $\neg Ib \rightarrow (Fab \leftrightarrow \forall x(x \varepsilon a \leftrightarrow (Ix \wedge (x \varepsilon b \vee \exists y(x \varepsilon y \wedge y \varepsilon b))))))$
- (c) $Fab \rightarrow Ma$
- (d) $Fab \rightarrow (Ia \leftrightarrow Ib)$
- (e) $\exists! x Fxb$

Proof. (a) Assume Mb . Either Ib or $\neg Ib$. Suppose Ib . In that case, $y \varepsilon^* b$ iff $y = b$, and so:

$$\begin{aligned} Fab &\leftrightarrow \forall x(x \varepsilon^* a \leftrightarrow \exists y(x \varepsilon^* y \wedge y = b)) \\ &\leftrightarrow \forall x(x \varepsilon^* a \leftrightarrow \exists y(x = y \wedge y = b)) \\ &\leftrightarrow \forall x(x \varepsilon^* a \leftrightarrow x = b) \end{aligned}$$

By Plurality, it cannot be the case that $\forall x(x \varepsilon a \leftrightarrow x = b)$. It follows that $x \varepsilon^* a$ iff $x = a$:

$$\begin{aligned} Fab &\leftrightarrow \forall x(x = a \leftrightarrow x = b) \\ &\leftrightarrow a = b \end{aligned}$$

Now suppose $\neg Ib$. In that case, $y \varepsilon^* b$ iff $y \varepsilon b$, and so:

$$\begin{aligned} Fab &\leftrightarrow \forall x(x \varepsilon^* a \leftrightarrow \exists y(x \varepsilon^* y \wedge y \varepsilon b)) \\ &\leftrightarrow \forall x(x \varepsilon^* a \leftrightarrow \exists y((Iy \wedge x = y) \vee x \varepsilon y) \wedge y \varepsilon b) \\ &\leftrightarrow \forall x(x \varepsilon^* a \leftrightarrow ((Ix \wedge x \varepsilon b) \vee \exists y(x \varepsilon y \wedge y \varepsilon b))) \end{aligned}$$

Since $Mb, \forall x(x \varepsilon b \rightarrow Ix)$. So:

$$Fab \leftrightarrow \forall x(x \varepsilon^* a \leftrightarrow x \varepsilon b)$$

By Plurality, b must include more than one thing, and so $\neg Ia$. This then implies that $x \varepsilon^* a$ iff $x \varepsilon a$, yielding:

$$Fab \leftrightarrow \forall x(x \varepsilon a \leftrightarrow x \varepsilon b)$$

So, by Extensionality:

$$Fab \leftrightarrow a = b$$

□

Proof. (b) Assume $\neg Ib$. In that case, $y \varepsilon^* b$ iff $y \varepsilon b$, and so:

$$\begin{aligned} Fab &\leftrightarrow \forall x(x \varepsilon^* a \leftrightarrow \exists y(x \varepsilon^* y \wedge y \varepsilon b)) \\ &\leftrightarrow \forall x(x \varepsilon^* a \leftrightarrow \exists y(((Iy \wedge x = y) \vee x \varepsilon y) \wedge y \varepsilon b)) \\ &\leftrightarrow \forall x(x \varepsilon^* a \leftrightarrow ((Ix \wedge x \varepsilon b) \vee \exists y(x \varepsilon y \wedge y \varepsilon b))) \end{aligned}$$

By Super-Plurality, $\forall x \forall y((x \varepsilon y \wedge y \varepsilon b) \rightarrow Ix)$. So:

$$\begin{aligned} Fab &\leftrightarrow \forall x(x \varepsilon^* a \leftrightarrow ((Ix \wedge x \varepsilon b) \vee \exists y(Ix \wedge x \varepsilon y \wedge y \varepsilon b))) \\ &\leftrightarrow \forall x(x \varepsilon^* a \leftrightarrow (Ix \wedge (x \varepsilon b \vee \exists y(x \varepsilon y \wedge y \varepsilon b)))) \end{aligned}$$

By Plurality, Fab then implies that $\neg Ia$, and thus that $x \varepsilon^* a$ iff $x \varepsilon a$:

$$Fab \leftrightarrow \forall x(x \varepsilon a \leftrightarrow (Ix \wedge (x \varepsilon b \vee \exists y(x \varepsilon y \wedge y \varepsilon b))))$$

□

Proof. (c) A straightforward corollary of (a) + (b). □

Proof. (d) A straightforward corollary of (a) + (b) + Plurality + Super-Plurality. □

Proof. (e) Either Ib or $\neg Ib$. If Ib , then $\exists! x Fxb$ is a trivial corollary of (a). If $\neg Ib$, then $\exists! x Fxb$ is a trivial corollary of (b) + Plurality + Comprehension + Extensionality. □

C Introducing parthood

Definition 4: Definitions of improper parthood, proper parthood, and overlap:

- (a) $a \leq b := \exists x(Fbx \wedge a \varepsilon^* x)$
- (b) $a < b := a \leq b \wedge a \neq b$
- (c) $a \circ b := \exists x(x \leq a \wedge x \leq b)$

We can now establish some important lemmas:

Lemma 5: A list of elementary facts about parthood and overlap:

- (a) $Ia \rightarrow (a \leq b \leftrightarrow (a \varepsilon^* b \wedge Mb))$
- (b) $\neg Ia \rightarrow (a \leq b \leftrightarrow (a \sqsubseteq b \wedge Mb))$
- (c) $Ma \leftrightarrow a \leq a$
- (d) $Ib \rightarrow (a \leq b \leftrightarrow a = b)$
- (e) $(Ia \wedge Mb) \rightarrow ((a \circ b \vee b \circ a) \leftrightarrow a \varepsilon^* b)$
- (f) $(Ma \vee Mb) \rightarrow (a = b \leftrightarrow \forall x(x \leq a \leftrightarrow x \leq b))$

Proof. (a) Assume Ia .

Left to Right: Assume $a \leq b$, with c as a witness to the existential: $Fbc \wedge a \varepsilon^* c$. If Ic , then $b = c$ by Lemma 3a, and so $a \varepsilon^* b \wedge Mb$. If $\neg Ic$, then $\neg Ib$ by Lemma 3d, and so $a \varepsilon b \wedge Mb$ by Lemmas 3b and c.

Right to Left: Assume $a \varepsilon^* b \wedge Mb$. By Lemma 3a, $Fbb \wedge a \varepsilon^* b$, and so $a \leq b$. □

Proof. (b) Assume $\neg Ia$.

Left to Right: Assume $a \leq b$, with c as a witness to the existential: $Fbc \wedge a \varepsilon^* c$. Mb follows immediately by Lemma 3c. Since $\neg Ia$, it follows that $a \varepsilon c$. By Super-Plurality, it follows that Sc and Pa . Lemma 3d then implies that $\neg Ib$. So, by Lemma 3b, we finally have $a \sqsubseteq b$.

Right to Left: Assume $a \sqsubseteq b \wedge Mb$. It immediately follows that Ma . So, since $\neg Ia$, there is some c s.t. $Ic \wedge c \varepsilon a$. Since $a \neq c$, Comprehension entails that there is some d s.t. $\forall x(x \varepsilon d \leftrightarrow (x = a \vee x = b \vee x = c))$. Clearly, $a \varepsilon d$, and by Lemma 3b, Fbd . So $a \leq b$. \square

Proof. (c) A trivial corollary of (a) + (b). \square

Proof. (d) Another trivial corollary of (a) + (b). \square

Proof. (e) A trivial corollary of (a) + (d). \square

Proof. (f) A trivial corollary of (a) + (b) + Extensionality. \square

D Fusion and parthood together

We have defined parthood in terms of fusion, but the standard approach is to go the other way around, and define fusion in terms of parthood. Because we have taken a non-standard route, we need to ensure that fusion and parthood are related in the right way. The key to demonstrating this is the following theorem:

Theorem 6: $Ma \rightarrow$

$$(Fab \leftrightarrow (\forall x(x \varepsilon^* b \rightarrow x \leq a) \wedge \forall x(x \leq a \rightarrow \exists y(y \varepsilon^* b \wedge x \circ y))))$$

If we assume Ib , then Theorem 6 becomes the following triviality (by Plurality + Lemmas 3 & 5):

$$Ma \rightarrow (a = b \leftrightarrow \forall x(x = b \leftrightarrow x = a))$$

So, in what follows, we will assume Ma and $\neg Ib$. Our aim is to prove this biconditional:

$$Fab \leftrightarrow (\forall x(x \varepsilon b \rightarrow x \leq a) \wedge \forall x(x \leq a \rightarrow \exists y(y \varepsilon b \wedge x \circ y)))$$

We can split our proof into four steps:

- (a) $Fab \rightarrow (c \varepsilon b \rightarrow c \leq a)$
- (b) $Fab \rightarrow (c \leq a \rightarrow \exists y(y \varepsilon b \wedge c \circ y))$
- (c) $\exists x(x \varepsilon^* a \wedge \neg \exists y(x \varepsilon^* y \wedge y \varepsilon b)) \rightarrow \exists x(x \leq a \wedge \neg \exists y(y \varepsilon b \wedge x \circ y))$
- (d) $\exists x(\neg x \varepsilon^* a \wedge \exists y(x \varepsilon^* y \wedge y \varepsilon b)) \rightarrow \exists x(x \varepsilon b \wedge \neg x \leq a)$

Steps (a) and (b) obviously establish the left-to-right reading of the biconditional; and given our assumptions, Steps (c) and (d) establish (the contraposition of) the right-to-left reading.

Proof. (a) Assume Fab . Suppose Ic and $c \varepsilon b$: Lemma 3b implies that $c \varepsilon b \rightarrow c \varepsilon a$; so, by Lemma 5a, $c \leq a$. Now suppose $\neg Ic$ and $c \varepsilon b$: by Super-Plurality, $\forall x(x \varepsilon c \rightarrow Ix)$; so, by Lemma 3b, $c \sqsubseteq a$, and hence $c \leq a$ by Lemma 5b. So, whether Ic or $\neg Ic$, $c \varepsilon b \rightarrow c \leq a$. \square

Proof. (b) Assume Fab . Suppose Ic and $c \varepsilon^* a$ (which is equivalent to $c \leq a$ by Lemma 5a): by the definition of fusion (Definition 2), $\exists y(c \varepsilon^* y \wedge y \varepsilon b)$, which is then equivalent to $\exists y(y \varepsilon b \wedge c \circ y)$ by Lemma 5e. Now suppose $\neg Ic$ and $c \sqsubseteq a$ (which is equivalent to $c \leq a$ by Lemma 5b): by the definition of fusion, $\exists z(Iz \wedge z \varepsilon c \wedge \exists y(y \varepsilon b \wedge z \varepsilon^* y))$; Super-Plurality + Lemma 5a then implies $\exists z(z \leq c \wedge \exists y(y \varepsilon b \wedge z \leq y))$, which is equivalent to $\exists y(y \varepsilon b \wedge c \circ y)$. So whether Ic or $\neg Ic$, $c \leq a \rightarrow \exists y(y \varepsilon b \wedge c \circ y)$. \square

Proof. (c) Assume the antecedent, with c as a witness:

$$c \varepsilon^* a \wedge \neg \exists y(c \varepsilon^* y \wedge y \varepsilon b)$$

Since Ma , it follows that Ic , and thus $c \leq a$ by Lemma 5a. Lemma 5e then implies that $\forall y(c \circ y \leftrightarrow c \varepsilon^* y)$. So we have:

$$c \leq a \wedge \neg \exists y(y \varepsilon b \wedge c \circ y)$$

\square

Proof. (d) Assume the antecedent, with c and d as witnesses:

$$\neg c \varepsilon^* a \wedge c \varepsilon^* d \wedge d \varepsilon b$$

$c \varepsilon^* d \wedge d \varepsilon b$ implies Ic and Id by Super-Plurality. It follows that $\neg c \leq a$ by Lemma 5a. If Id , then $c \varepsilon b$, and so $\exists x(x \varepsilon b \wedge \neg x \leq a)$. If $\neg Id$, then $c \varepsilon d$, and so $\neg d \sqsubseteq a$, which implies $\neg d \leq a$ by Lemma 5b; so again, we have $\exists x(x \varepsilon b \wedge \neg x \leq a)$. \square

Theorem 6 is enough to establish that fusion and parthood relate as they should. Here is one way to see this. Against the backdrop of a singular first-order logic, fusion is expressed as $F^*(x, \phi)$, which is stipulated to obey these laws (where y does not appear in ϕ):

$$(\exists x \phi \rightarrow \exists x F^*(x, \phi)) \wedge (F^*(a, \phi) \leftrightarrow (\forall x(\phi \rightarrow x \leq a) \wedge \forall x(x \leq a \rightarrow \exists y(\phi_{y/x} \wedge z \circ a))))$$

We can now define $F^*(x, \phi)$ as follows (where y and z do not appear in ϕ):

Definition 7: $F^*(x, \phi) := \exists y((\forall z(\phi_{z/x} \leftrightarrow y = z) \vee \forall z(\phi_{z/x} \leftrightarrow z \varepsilon y)) \wedge Fxy)$

It would then be straightforward to prove the desired result (where y does not appear in ϕ , and the quantifiers have been appropriately restricted to M):

Theorem 8: $(\exists x(Mx \wedge \phi) \rightarrow \exists x F^*(x, \phi)) \wedge (F^*(a, \phi) \leftrightarrow (\forall x(Mx \rightarrow (\phi \rightarrow x \leq a)) \wedge \forall x(x \leq a \rightarrow \exists y(\phi_{y/x} \wedge a \circ y))))$

Proof. A trivial corollary of Theorem 6 + Comprehension. \square

E Classical Mereology

Here is an axiomatisation of Classical Mereology (CM):

Transitivity: $\forall x \forall y \forall z ((x \leq y \wedge y \leq z) \rightarrow x \leq z)$

Universalism: $\forall x \exists y Fyx$

Weak Supplementation: $\forall x \forall y (x < y \rightarrow \exists z (z \leq y \wedge \neg x \circ z))$

Transitivity is trivial, and Universalism has already been proved (Lemma 3e). Here is a proof of Weak Supplementation:

Proof. Assume $a < b$. It immediately follows that Ma and Mb . Suppose Ia : by Plurality, there is some c s.t. $c \varepsilon b \wedge c \neq a$; since Mb , it follows that Ic , and so, by Lemmas 5a & 5e, $c \leq b \wedge \neg a \circ c$. Now suppose $\neg Ib$: by Lemma 5b, $a \sqsubset b$, and so there is some c s.t. $c \varepsilon b \wedge \neg c \varepsilon a$; since Mb , it follows that Ic , and so, again by Lemmas 5a & 5e, $c \leq b \wedge \neg a \circ c$. So, no matter whether Ia or $\neg Ia$, $\exists z (z \leq b \wedge \neg a \circ z)$. \square

Acknowledgments

We would like to thank Tim Button, Dave Ingram, Chris Jay, Owen Hulatt, audiences at York and Warsaw, and our anonymous reviewers. Simon Thunder was employed on the postdoctoral research programme at UNAM, and he gratefully acknowledges their support. The details of that programme are: UNAM, Programa de Becas Posdoctorales en la UNAM, Becario del Instituto de Investigaciones Filosóficas, asesorado por el doctor Ricardo Mena.

References

- Barker, Jonathan (2021). ‘Grounding and the myth of ontological innocence’. *Australasian Journal of Philosophy* 99, pp.303–318.
- Baxter, Donald L. M. (1988). ‘Many-one identity’. *Philosophical Papers* 17, pp.193–216.
- (2014). ‘Identity, discernibility, and composition’. In: Cotnoir and Baxter 2014, pp.244–254.
- Bennett, Karen (2015). ‘“Perfectly understood, unproblematic, and certain”: Lewis on mereology’. In: *A Companion to David Lewis*. Ed. by Barry Loewer and Jonathan Schaffer. Hoboken: Wiley Blackwell, pp.250–261.
- Bennett, Karen and Dean W. Zimmerman, eds. (2013). *Oxford Studies in Metaphysics*. Vol. 8. Oxford: Oxford University Press.
- Berry, Sharon (2022). *A Logical Foundation for Potentialist Set Theory*. Cambridge: Cambridge University Press.
- Bøhn, Einar Duenger (2014). ‘Unrestricted Composition as Identity’. In: Cotnoir and Baxter 2014, pp.143–166.
- (2016). ‘Composition as Identity and Plural Cantor’s Theorem’. *Logic and Logical Philosophy* 25, pp.411–428.
- Boolos, George (1984). ‘To be is to be the value of a variable (or to be some values of some variables)’. *The Journal of Philosophy* 72, pp.430–449.
- Bricker, Phillip (2016). ‘Composition as a Kind of Identity’. *Inquiry* 59, pp.264–294.
- (2021). ‘Composition as Identity, Leibniz’s Law, and slice-sensitive emergent properties’. *Synthese* 198 (Suppl 18), S4389–S4409.

- Button, Tim and Robert Trueman (2022). 'Against cumulative type theory'. *The Review of Symbolic Logic* 15, pp.907–949.
- Calosi, Claudio (2016). 'Composition as Identity and mereological nihilism'. *The Philosophical Quarterly* 66, pp.219–235.
- Cameron, Ross P. (2007). 'The contingency of composition'. *Philosophical Studies* 136, pp.99–121.
- (2012). 'Composition as Identity doesn't settle the Special Composition Question'. *Philosophy and Phenomenological Research* 84, pp.531–554.
- (2014). 'Parts generate the whole, but they are not identical to it'. In: Cotnoir and Baxter 2014, pp.90–108.
- Carnap, Rudolf (1945). 'The two concepts of probability: The problem of probability'. *Philosophy and Phenomenological Research* 5, pp.513–532.
- (1950). *Logical Foundations of Probability*. Chicago, IL: University of Chicago Press.
- Carrara, Massimiliano and Giorgio Lando (2016). 'Composition, indiscernibility and coreferentiality'. *Erkenntnis* 81, pp.119–142.
- Cotnoir, Aaron J. (2013). 'Composition as general identity'. In: Bennett and Zimmerman 2013, pp.294–322.
- (2014). 'Composition as identity: Framing the debate'. In: Cotnoir and Baxter 2014, pp.3–23.
- Cotnoir, Aaron J. and Donald L. M. Baxter, eds. (2014). *Composition as Identity*. Oxford: Oxford University Press.
- Cotnoir, Aaron J. and Achille C. Varzi (2021). *Mereology*. Oxford: Oxford University Press.
- Fine, Kit (2006). 'Relatively unrestricted quantification'. In: Rayo and Uzquiano 2006, pp.20–44.
- Florio, Salvatore and Øystein Linnebo (2016). 'On the innocence and determinacy of plural quantification'. *Noûs* 50, pp.565–583.
- (2021). *The One and the Many: A Philosophical Study of Plural Logic*. Oxford: Oxford University Press.
- Grimau, Berta (2021). 'In defence of Higher-Level Plural Logic: Drawing conclusions from natural language'. *Synthese* 198, pp.5253–5280.
- Harte, Verity (2002). *Plato on Parts and Wholes: The Metaphysics of Structure*. Oxford: Oxford University Press.
- Hawley, Katherine (2013). 'Cut the pie any way you like? Cotnoir on general identity'. In: Bennett and Zimmerman 2013, pp.323–330.
- (2014). 'Ontological innocence'. In: Cotnoir and Baxter 2014, pp.70–89.
- Hewitt, Simon Thomas (2012). 'Modalising plurals'. *Journal of Philosophical Logic* 41, pp.853–875.
- Hodva, Paul (2009). 'What is Classical Mereology?' *Journal of Philosophical Logic* 38, pp.55–82.
- Hossack, Keith (2000). 'Plurals and complexes'. *British Journal for the Philosophy of Science* 51, pp.411–443.
- Inwagen, Peter van (1994). 'Composition as identity'. *Philosophical Perspectives* 8, pp.207–220.
- Lewis, David (1991). *Parts of Classes*. Oxford: Basil Blackwell.
- Linnebo, Øystein (2010). 'Pluralities and sets'. *The Journal of Philosophy* 107, pp.144–164.
- (2018a). 'Dummett on indefinite extensibility'. *Philosophical Issues* 28, pp.196–220.
- (2018b). *Thin Objects*. Oxford: Oxford University Press.
- Linnebo, Øystein and David Nicolas (2008). 'Superplurals in English'. *Analysis* 68, pp.186–197.
- Linnebo, Øystein and Agustín Rayo (2012). 'Hierarchies ontological and ideological'. *Mind* 121, pp.269–308.
- Loss, Roberto (2018). 'A sudden collapse to nihilism'. *Philosophical Quarterly* 68, pp.370–375.

- Loss, Roberto (2021a). 'Composition, identity and plural ontology'. *Synthese* 198, pp.9193–9210.
- (2021b). 'On Atomic Composition as Identity'. *Synthese* 198 (Suppl 18), S4519–S4542.
- (2022). 'Composition as Identity and the innocence of mereology'. *Philosophy and Phenomenological Research* 105, pp.128–143.
- McDaniel, Kris (2010). 'Composition as Identity does not entail universalism'. *Erkenntnis* 73, pp.97–100.
- Merricks, Trenton (2005). 'Composition and vagueness'. *Mind* 114, pp.615–637.
- Oliver, Alex and Timothy Smiley (2016). *Plural Logic*. 2nd ed. Oxford: Oxford University Press.
- Payton, Jonathan D. (2021a). 'Composition as Identity, now with all the pluralities you could want'. *Synthese* 199, pp.8047–8068.
- (2021b). 'How to identify wholes with their parts'. *Synthese* 198 (Suppl 18), S4571–S4593.
- (2022). 'Composition and plethological innocence'. *Analysis* 82, pp.67–74.
- (2025). 'From singular to plural . . . and beyond?' *Philosophy and Phenomenological Research* 110, pp.983–1012.
- Quine, Willard Van Orman (1948). 'On what there is'. In: Quine 1980, pp.1–19.
- (1980). *From a Logical Point of View*. 2nd ed. Cambridge, MA: Harvard University Press.
- Rayo, Agustín (2006). 'Beyond plurals'. In: Rayo and Uzquiano 2006, pp.220–254.
- (2007). 'Ontological commitment'. *Philosophy Compass* 2, pp.428–444.
- Rayo, Agustín and Gabriel Uzquiano, eds. (2006). *Absolute Generality*. Oxford: Oxford University Press.
- Russell, Jeffrey Sanford (2017). 'Composition as abstraction'. *The Journal of Philosophy* 114, pp.453–470.
- Saucedo, Raul (2022). 'Ontological collectivism'. *Philosophical Perspectives* 36, pp.233–269.
- (2025). 'Identity universalism and ontological explosion'. *Synthese* 205, pp.1–25.
- Shapiro, Stewart and Crispin Wright (2006). 'All things indefinitely extensible'. In: Rayo and Uzquiano 2006, pp.255–304.
- Sider, Theodore (2007). 'Parthood'. *The Philosophical Review* 116, pp.51–91.
- (2011). *Writing the Book of the World*. Oxford: Clarendon Press.
- (2013). 'Against parthood'. In: Bennett and Zimmerman 2013, pp.237–293.
- (2014). 'Consequences of collapse'. In: Cotnoir and Baxter 2014, pp.211–222.
- Studd, James P. (2019). *Everything, More or Less: A Defence of Generality Relativism*. Oxford: Oxford University Press.
- Thunder, Simon (2023). 'Composite objects are mere manys'. In: *Oxford Studies in Metaphysics*. Ed. by Karen Bennett and Dean W. Zimmerman. Vol. 13. Oxford: Oxford University Press, pp.20–50.
- Trueman, Robert (2014). 'A dilemma for neo-Fregeanism'. *Philosophia Mathematica* 22, pp.361–379.
- (2021). *Properties and Propositions: The Metaphysics of Higher-Order Logic*. Cambridge: Cambridge University Press.
- Turner, Jason (2014). 'Donald Baxter's Composition as Identity'. In: Cotnoir and Baxter 2014, pp.225–243.
- Werner, Jonas (2022). 'Irreducibly collective existence and bottomless nihilism'. *Synthese* 200, pp.1–16.
- Williams, J. Robert G. (2006). 'Illusions of gunk'. *Philosophical Perspectives* 20, pp.493–513.
- Yablo, Stephen (2006). 'Circularity and paradox'. In: *Self-Reference*. Stanford, CA: CSLI Publications, pp.139–157.
- Yi, Byeong-uk (1999). 'Is mereology ontologically innocent?' *Philosophical Studies* 93, pp.141–160.
- (2014). 'Is there a plural object?' In: Cotnoir and Baxter 2014, pp.169–191.

Yi, Byeong-uk (2021). 'Is composition identity?' *Synthese* 198 (Suppl 18), S4467–S4501.