Wholes and Parts: The Limits of Composition[1]

D. H. Mellor

Cambridge University
Faculty of Philosophy
Sidgwick Avenue
Cambridge
UK, CB1 1JS
dhm11@cam.ac.uk

Abstract:
The paper argues that very different part-whole relations hold between different kinds of entities. While these relations share most of their formal properties, they need not share all of them. Nor need other mereological principles be true of all kinds of part-whole pairs. In particular, it is argued that the principle of unrestricted composition, that any two or more entities have a mereological sum, while true of sets and propositions, is false of things and events.

1. Part-whole relations

Entities of many different kinds have been taken to be wholes or parts, as in the following list adapted from Armstrong (1978:36):

> The proposition that *p* is a part of the proposition that *p&q*;
> The property *F* is a part of the property *F&G*.
> The set of women is a part of the set of human beings;
> New South Wales is a part of Australia;
> The Terror was a part of the French Revolution.

What enables entities of all these very different kinds to be related to each other as parts to wholes? The obvious, if hardly enlightening, answer is that we can apply a concept of *inclusion* to all of them: Australia includes New South Wales, The French Revolution includes The Terror, the set of human beings includes that of women, and so on.

The wide range of these examples suggests that our concept of inclusion is *formal* rather than *material*. In other words, we apply the concept to pairs of entities linked by relations with the formal properties of inclusion, however much those relations may differ in other ways. Take the relation that makes a proposition *p* part of the proposition *p&q*. This relation seems to be *entailment*, since what makes *p&q* include *p* seems to be that *p&q* entails *p*. (This captures the idea that the conclusion of a valid argument is in some serious sense *contained* in its premises.)

Similarly, perhaps, with properties, since if anything has the property *F* and the

---

property *G*, this entails that it has the property *F*. Alternatively, if we identify properties with the sets of all the possible particulars that have those properties (Lewis, 1986 ch. 1.5), we might take *set inclusion* to be what makes one property part of another, just as it is what makes the set of women part of the set of all human beings. That, however, reverses the apparent part-whole relation in this case; for since the set of all *F*-particulars includes the set of all *F&G*-particulars, the set-inclusion relation makes the property *F&G* look like part of the property *F* rather than the other way round. The same thing happens with propositions, if we identify these with the sets of possible worlds in which they are true (Lewis, 1986:53-55). Here too, since the set of worlds where *p* is true includes the set in which *p&q* is true, set inclusion makes *p&q* look like a part of *p*.

So, back to entailment then; or perhaps not. The trouble with entailment is that, in classical logic, every proposition entails all *necessary* propositions and is entailed by all *impossible* ones. This seems to make every necessary proposition a part of all propositions, and all propositions parts of every impossible proposition, thus making the proposition that 2+2=4 part of the proposition that the earth is round, and this proposition in turn part of the proposition that 2+2=5. However, we can avoid this unattractive implication by applying our part-whole concept only to *contingent* propositions, where the entailment relation may more plausibly be what makes one such proposition part of another.

There are other arguments for and against entailment and set inclusion as part-whole relations, which we need not go into here. The point here is simply that one of them may be a more credible part-whole relation for propositions and the other for properties. And that is the point: the relation that makes one entity part of another may vary from one kind of entity to another.

The other examples in our list reinforce that point. Take the part-whole relation between New South Wales and Australia. We might take this to be set inclusion if we identified these geographical entities with sets of points on the earth's surface. But which sets? Obviously the sets of all and only those points that lie within the spatial boundaries of New South Wales and Australia. But then the relation that determines which pairs of regions of the earth's surface are related as parts to wholes will be *spatial* rather than *set* inclusion.

Similarly with the Terror and the French Revolution, except that now we need *temporal* as well as spatial inclusion. And even that is not enough, since not everything that happened within France during the French Revolution is a *part* of that revolution. For example, few if any births, movements or deaths of wild animals within France at that time were parts of it; and if some human deaths and movements, and perhaps even a few human births, were parts of it, most were not.

What else then, besides spatiotemporal inclusion, does it take to make an event part of the French Revolution? It is hard to say, precisely. But here it will do to say that the event has to be what I shall call a *working part* of the revolution. In other words, no event is a part of the French Revolution unless some relevant features of the revolution 'as a whole' depend on that event *causally*, as they certainly did on The Terror and some of its working parts.

This condition, that parts of events must be working parts, is not peculiar to revolutions. Take a three-course dinner for several people. This event – the dinner – looks like a whole, with its courses as temporal parts, which in turn have their consumption by the diners as spatial parts. Again, not every event within the dinner's spatial and

temporal confines is a part of it: for example, unnoticed movements of insects across the table are not. Here too the extra condition seems to be causal. Only events that make some causal contribution to relevant properties of the dinner as a whole will be parts of it.

As for events, so for *continuants*, or *things*, as I shall call them for short, including not only inanimate things, both natural and man-made, but plants, animals, people, and all the other entities that provide the most obvious instances of part-whole pairs. Take artefacts, such as cars. Cars are clearly material wholes, whose parts include their engines, gear boxes, wheels, doors, etc.; most of which are themselves wholes with parts that include the constituents from which they in turn are assembled.

This is not of course to say that all an artefact's parts are things it is assembled from. Artefacts like statues made not by being assembled, but by being carved out of solid blocks of wood or stone, still *have* parts. But equally, not all entities within the spatial boundaries of an artefact, assembled or not, are parts of it. Airline passengers are not parts of the planes they fly in, any more than photons passing through a window are, while they do so, parts of the window.

With artefacts, in short, as with events, what we mean by a 'part' is a working part, something on which a whole thing depends causally for some properties relevant to the kind of thing it is – a car, a statue, an aircraft, a window. And as for artefacts, so for natural things like planets and rocks, as well as plants, animals in general, and people in particular. We also take all their parts to be working parts (cf. Simons, 1987: 107), thereby excluding, for example, radio waves that pass through them.

This having been said, let us return to our initial list of part-whole pairs: the proposition p and the conjunction $p\&q$; the property F and the property $F\&G$; the sets of women and of human beings; New South Wales and Australia; and the Terror and the French Revolution; to which, for completeness, we may add pairs of things, such as our cells and the human bodies they are parts of.

Notice again how *heterogeneous* this list is. The entities that it says are related as parts to wholes are pairs, respectively, of propositions, properties, sets, geographical regions, events and things. But equally striking, given this heterogeneity, is the *homogeneity* of each pair. In none of them is the whole different in kind from the part. Properties and propositions are not paired with each other, geographical regions are not paired with sets, things are not paired with events, and so on.

There are admittedly exceptions to this rule. Take the proposition $\langle Fa \rangle$ that a particular thing a has the property F. This proposition looks like a whole with a and F as parts, even though a and F are not propositions and so cannot be made parts of $\langle Fa \rangle$ just by the entailment relation that makes one proposition part of another. Still, these and other apparent exceptions do not alter the fact that, as a rule, part-whole relations relate entities of the same kind. The best way to explain both the rule and the exceptions to it is to allow a multiplicity of part-whole relations, from spatiotemporal inclusion plus causal dependence to set inclusion, entailment, and whatever relation gives propositions their non-propositional parts.

But of all these examples it is events and things that make the strongest case against a single part-whole relation. As we have seen, for one event or thing to be a part of another, at least two necessary conditions must be met: spatiotemporal inclusion, and some kind of causal dependence of the whole on the part. And whatever we take causation to be, these two conditions are independent. Neither is reducible either to the

other or to any of the quite different relations that make one proposition, property or set part of another.

2. The formal properties of part-whole relations

The multiplicity of part-whole relations may or may not show that we have more than one part-whole *concept*, as some philosophers think (Simons, 1987:128). That seems to me a moot but unimportant point, just as it is, for example, whether the material differences between light, sound and surface waves show that we have more than one concept of a *wave*. What the existence of different part-whole relations certainly does show is that, as noted earlier, our part-whole concept (or family of concepts) is a formal rather than a material one, as indeed our wave concept (or family of concepts) is.

For in the wave case, all something has to do to be a wave, whatever it is made of, is to satisfy a wave *equation*. That formal fact about light, sound and surface waves is what makes us call them waves. Our doing so does not conflict with the fact that light waves are material features of light (i.e. of electromagnetic radiation), sound waves are quite different features of air or other fluids, and surface waves are equally different features of water or other surfaces. On the contrary: being features of quite different things is precisely what allows us, if we wish, to apply a single wave concept to all these different features, by stopping the features from conflicting. (It is what, for example, stops sound waves also being light waves of a different frequency.) Similarly with parts and wholes. What stops different part-whole relations conflicting, thereby enabling us, if we wish, to bring them all under a single formal part-whole concept, is the fact that they relate different entities.

This is why mereology is usually, and correctly, regarded as 'the *formal* theory of part, whole and related concepts' (Simons 1987: 1 – my italics). But this way of reading it exposes mereologists to the occupational hazard of trying to decide *a priori* what the formal properties of part-whole relations are. For while, obviously, different part-whole relations need to share many formal properties if they are to fall under a single formal concept, they need not share *all* of them.

Take the question of whether part-whole relations are *transitive*, *intransitive* or neither, i.e. whether if A is a part of B and B is a part of C, it follows that A *is* a part of C, is *not* a part of C, or may or may not be a part of C. In fact most mereologists assume that part-whole relations *are* transitive (Simons, 1987 ch.3.2.2), because most of them – entailment, set inclusion and spatiotemporal inclusion – obviously are. If A entails B and B entails C, then A entails C, and similarly if A includes B and B includes C, whether A, B and C are sets, or regions of space, time or space-time.

However, whether *causal dependence* is transitive is less obvious, especially where the dependence is not deterministic (Lewis, 1973:167; Mellor, 2003:228). So if events or things with parts must be causally dependent on those parts, then whether transitivity holds here is a question for our theory of causation, a question whose answer will then affect what parts events and things have. Thus, if causal dependence *is* transitive, my bodily parts will include all *their* parts, and so on, down to the smallest sub-atomic parts of every atom in my body. But if it is *not* transitive, then the parts of my body will only include those of their parts that affect my body as a whole.

This shows why we must derive the formal properties of our part-whole concept from those of part-whole relations, not the other way round. We cannot derive them from *a priori* intuitions about parts and wholes in general. Doing that is like trying to decide *a priori* whether waves are longitudinal or transverse, i.e. whether they oscil-

late in the direction the wave moves in, or at right angles to it: a hopeless project, since different waves oscillate differently, sound waves being longitudinal and light and surface waves transverse.

Now no one, so far as I know, ever did try to decide *a priori* which way waves oscillate. But many nineteenth century scientists did think they knew *a priori* that all waves need a material medium, when the fact is that some do and some do not. Sound waves, for example, do need such a medium: air or water, or some other material fluid; just as surface waves need a boundary, like that between air and a body of water that is held down by gravity. Light, on the other hand, needs no such medium – not counting electromagnetic fields, which will not do what was wanted, namely provide something for the speed of light to be relative to, as the speed of sound and surface waves is relative to the air or water in which they are waves. Hence the late nineteenth century search for a material ether, which we now know does not exist.

3. Unrestricted composition?

Bearing this cautionary tale in mind, let us return to mereology and, in particular, to the *principle of unrestricted composition*, or PUC for short, which says that all pluralities of entities compose wholes that have those entities as parts. PUC is widely accepted (Rea, 1998), although not universally so (Chisholm, 1976:219; Rescher, 1955). Is PUC true? The answer, as with transitivity, I shall now argue depends on the kind of entities we are talking about: PUC is true of some kinds and false of others.

First, we must note the insignificance of the fact that, in mereology, wholes are called the *sums* of their parts. Since this is true by mere definition, being simply what the term 'sum' means here, and so does not entail that any entities either are or have sums (other than themselves), it does not entail that composition is unrestricted.

Next, as with transitivity, we must note that there are several kinds of entities of which PUC obviously is true; otherwise, of course, no one would have assumed it. For example, it is obviously true of *sets*. But this is only because it is obvious what we should take the sum of two or more sets, such as the set of women and the set of men, to be: namely their *union*, in this case the set of which all and only women and men are members. This is the assumption that makes PUC true of sets, by making PUC follow from the fact that any number of sets *has* a union, which is also a set. And however uncontentious that assumption is, it takes more than mereology to justify it, because it is an assumption about sets, not just about sums.

Similarly for *propositions*, where the relevant and equally uncontentious assumption is that the sum of any two or more propositions is their *conjunction*. This, and the fact that any number of propositions *have* a conjunction, which is also a proposition, is what makes PUC true of propositions.

Properties are less obvious, since it is a moot point whether the only credible candidates for being sums of properties, namely *conjunctive* properties like *F&G*, exist. Many philosophers think they do; often because, as we noted in Section 1, properties can be identified with the sets of all their possible instances. Since I reject this identification (I think properties are universals), I can, and in fact do, deny that there are any such entities as conjunctive properties (Mellor 1997: 264). If I am right, then there is *no* non-trivial composition of properties: properties have no properties other than themselves as parts. However, I need not argue this point here since, as with transitivity, all that matters is that, whether PUC is true of properties depends on a thesis about properties that cannot be established just by a general theory of parts and wholes.

How about regions of *space*, *time* or *spacetime*? If we identify such regions with the sets of the spatiotemporal points they contain, will that not make their composition as unrestricted as that of sets in general? Not necessarily. Take two disconnected lines, say one in London and one in Cape Town, or two disconnected time intervals, say 2005 and 2007. Do such lines and time intervals have sums? If they do, their sums are certainly not connected, and I see no non-question-begging reason to suppose that they exist, since theories of space, time and spacetime need only postulate internally connected regions.

Finally, and most importantly, is the composition of *things* and *events* unrestricted? Is there, for example, a sum of the Pope and the Kremlin, or of the Big Bang and the Boer War (and perhaps of their parts, if their part-whole relations are transitive)? Not on the face of it. Why then do so many mereologists believe in the unrestricted composition of things and events?

First, there is a formal reason for accepting PUC, namely that it is essential to a formal theory, that of classical extensional mereology, which provides a particularly powerful and elegant model of parts and wholes in general (Simons, 1987 ch.1.5). Models are however one thing, reality is another. To cross Siberia by train you need to know about the actual Trans-Siberian railway, not just about models of it. Similarly, the formal power and elegance of Euclidean geometry does not show that it is true of actual geodesics, and the success of the theory of general relativity shows that it is not. And so it is with mereology. The fact that a formally attractive model of parts and wholes needs PUC does not show that PUC is true of actual things and events. We need more than formal reasons to reject the hordes of apparent counter-examples.

And there are some such reasons. One is the utility in semantics of having singular terms to refer to arbitrary collections of entities; terms which expressions like 'the sum of ...' , 'the set of ...' (and of course 'the collection of ...') can be used to form. However, the merely semantic utility of singular terms like 'the set of Xs' or 'the sum of Xs' does not begin to show that in reality, besides the Xs that these terms provide useful ways of talking about collectively, there are any such entities as sets or sums of those Xs (Simons, 1987 ch.4.4).

There is, moreover, a crucial semantic difference between sets and sums which makes set theory more autonomous than mereology can be. The difference is that whereas sets are of a different semantic kind from whatever they are sets of, sums are not (Williams, 1953:117-18). Sets of events, for example, unlike sums of events, are not events. This means that, while we can debate the existence of sets in general, we can hardly debate that of sums in general, since what sums there are depends, as we have seen, on what they are sums of. Thus, since for example, sums of propositions are also propositions, no one who believes in propositions can very well deny the existence of *sums* of them, although they may well deny the existence of *sets* of them. Similarly, since the sums of events are events, and some events certainly have others as parts, no one who believes in events can very well deny that there *are* sums of events. And the same goes for things.

Still, this does nothing to show that *any* two or more things or events have a sum, i.e. that PUC is true of these kinds of entity. The claim that it is needs a more than formal or semantic theoretical utility to justify dismissing all our apparent counter-examples as mere 'folk ontology'. But again the question is: theoretical utility for what? – a question to which again the answer cannot be mereology in general, but can only be the theory of this application of it, namely the ontology of things and events. And to

that theory PUC is *not* useful. On the contrary, assuming that any two or more things or events have a sum, causes far more theoretical trouble than it saves, as an otherwise admirable article by David Lewis (1993) unintentionally shows.

In his paper, Lewis tackles problems posed for the mereology of things by the fact that it is sometimes a vague matter whether something is a part of something else.[2] He does this, following Geach (1967), by imagining a cat, Tibbles, many of whose loose hairs are borderline candidates for being among its parts. Now consider two sums of Tibbles's parts, which differ only in that one sum includes one such hair and the other does not. Each sum is an equally good candidate for being Tibbles, and yet they cannot both be Tibbles, since no hair can both be and not be a part of one and the same cat. Similarly for all the other borderline candidates for being parts of Tibbles, which gives us thousands of sums, all equally good but incompatible candidates for being that cat.

What then should we say, given that each of these myriad sums of cat-parts clearly *is* a cat, thus giving us thousands of actual cats where we thought there was only one? Lewis canvasses several possible solutions, with the one he favours postulating a relation of *partial identity* between all these overlapping cats.

We, however, can do far better by simply denying the principle of unrestricted composition, which generates Lewis's problem in the first place by entailing the existence of sums of any number of Tibbles's parts. Once we deny PUC, we can say that there is only the *one* cat, Tibbles; a cat that has many parts, including many of the hairs on its surface. It may indeed be a vague matter how many of those hairs *are* parts of Tibbles, just as it is a vague matter how many hairs Tibbles would have to lose to become bald. But this vagueness need only be semantic: it need not be ontological.

This is not of course to deny that the vagueness of predicates like 'is bald' poses serious problems, to which there is indeed no generally agreed solution (see Keefe and Smith 1996). My own view is that the least bad solution is the epistemic one advocated by Williamson (1994). But this is not a claim I need to argue here. All that matters here that, whatever the best solution, it will apply equally well to the vagueness of predicates like 'is a part of Tibbles'. For however unattractive some consequences of rival views of vagueness may be, none is as absurd as the mad multiplication of cats and other entities produced by applying PUC to things.

This absurdity is far worse than Lewis admits, as we can see by reflecting that if PUC is true of things in general, it will be true not only of all parts of Tibbles, but of all parts of Tibbles's hairs, each of which will also have many borderline candidates for being one of *its* parts, such as the dead or dying cells on its surface. This will then turn each apparently single hair into thousands, thus multiplying a thousand-fold Lewis's already ludicrously large number of cats. Nor will that be the end of it, since every *cell* of every hair will also have borderline candidates for being parts of *it*, such as molecules loosely attached to its surface. And so on and so on, down to quantum ensembles of fundamental particles that have no precise positions, but only a probability distribution over possible positions that extend to infinity, thus violating another basic mereological assumption about things, namely that the parts of things must be *inside* them.

In short, not only is PUC not a useful assumption in a mereological theory of the on-

---

2   Lewis does not explicitly discuss the problems that mereological vagueness poses for an ontology of events. But as these are essentially the same as those he does discuss, I too shall confine myself to things in the following discussion, leaving readers to apply the argument to events.

tology of things, it may well be fatal, since it makes that theory incompatible with modern microphysics. For while modern microphysics is indeed neither unproblematic nor complete, there is overwhelming evidence for the probabilistic features of it that conflict with the unrestricted composition of things. Given this conflict, the problems that vagueness makes PUC generate, the lack of any other theoretical argument for PUC, and its many apparent counter-examples, I conclude that PUC is false. The mereological composition of things is *not* unrestricted.

References

Armstrong, D. M. 1978. *A Theory of Universals*, Cambridge: Cambridge University Press.

Chisholm, R. M. 1976. *Person and Object: A Metaphysical Study*, London: Allen & Unwin.

Geach, P. T. 1967. 'Identity', in his *Logic Matters*. Oxford: Blackwell (1972), 238-47.

Keefe, R. and Smith, P., eds. 1996. *Vagueness: A Reader*. Cambridge, Mass.: MIT Press.

Lewis, D. K. 1973. 'Causation', in his *Philosophical Papers Volume II*. Oxford: Oxford University Press (1986), 159-72.

— 1986. *On the Plurality of Worlds*, Oxford: Blackwell.

— 1993. 'Many, but Almost One', in his *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press (1999), 164-82.

Mellor, D. H. 1997. 'Properties and Predicates', in *Properties*, ed. D. H. Mellor and Alex Oliver, Oxford: Oxford University Press, 255-67.

— 2003. '*Real Metaphysics*: Replies', in *Real Metaphysics*, ed. Hallvard Lillehammer and Gonzalo Rodriguez-Pereyra, London: Routledge, 212-38.

Rea, M.-C. 1998. 'In Defense of Mereological Universalism', *Philosophy and Phenomenological Research*, 63(2) 347-60.

Rescher, N. 1955. 'Axioms for the Part Relation', *Philosophical Studies* 6, 8-11.

Simons, P. 1987. *Parts: A Study in Ontology*, Oxford: Clarendon Press.

Williams, D. C. 1953. 'On the Elements of Being: I', in *Properties*, ed. D. H. Mellor and Alex Oliver, Oxford: Oxford University Press, 112-24.

Williamson, T. 1994. *Vagueness*, London: Routledge.