

DRAFT ONLY: NOT TO BE COPIED, QUOTED OR PUBLISHED

Wholes and Parts: The Limits of Composition¹

© D. H. Mellor 2006

1. Entities of many different kinds have been taken to be wholes or parts. Here, for example, is most of the list David Armstrong gives in *A Theory of Universals* (p.36):
2. 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 class of women is a part of the class of human beings;
New South Wales is a part of Australia;
The Terror was a part of the French Revolution;
3. What enables entities of all these very different kinds to be related to each other as parts to wholes? The obvious – if not very enlightening – answer is that we can apply a concept of *inclusion* to all of them.
4. Thus Australia *includes* New South Wales, The French Revolution *includes* The Terror, the class of human beings *includes* that of women, and so on.
5. The wide range of these examples suggests that this concept of inclusion is *formal* rather than *material*. By this I mean that we can apply the concept to pairs of entities linked by any relation with the formal properties of *inclusion*, however much that relation may differ in other ways from other such relations.
6. Take the relation that makes a proposition p part of the proposition $p \& q$. That relation seems to be *entailment*, since what makes $p \& q$ include p seems to be that $p \& q$ entails p . (This, for example, captures the idea that the conclusion of a valid argument is in some serious sense *contained* in its premises.)
7. Similarly, we might think, with properties: since if anything has both the property F and the property G , this *entails* that it has the property F .

¹ Earlier versions of this paper were discussed at a workshop at St Andrews University on 11/6/2005 and at a seminar at the Central European University on 11/10/2005. In revising the paper I have been greatly assisted by comments made in those discussions.

8. On the other hand, if we identify properties with the sets of all the possible particulars that have those properties, we might think that what makes the property F part of the property $F\&G$ is the relation of *set inclusion*.
9. For then what makes F part of $F\&G$ is that the set of all possible entities that are both F and G *includes* the set of all possible entities that are F ; just as what makes the class of women part of the class of all human beings includes that of all women.
10. So perhaps set-inclusion rather than entailment is the basic part–whole relation between properties. For even if properties *aren't* the sets of their possible instances, what makes being $F\&G$ entail being F is the fact that anything that's both F and G must be a member of the set of all F s.
11. But then might we not take set inclusion to be the part–whole relation for *propositions* too, which we could if we individuate propositions by the sets of possible worlds in which they're true?
12. That view, however, has a price, namely that, as the set of worlds where p is true includes the set where $p\&q$ is true, it makes $p\&q$ look like a part of p rather than the other way round.
13. So, back to entailment then; or perhaps not. The trouble here is the fact that, in classical logic, every proposition entails all *necessary* propositions and is entailed by all *impossible* ones, which seems to make every necessary proposition a part of all propositions, and all propositions parts of every impossible proposition.
14. And do we really want to say that the necessary proposition that $2+2=4$ is a *part* of the contingent proposition that the earth is round, and that this proposition in turn is a part of the impossible proposition that $2+2=5$?
15. Well, probably not; but then we can always avoid that consequence by only applying our part–whole concept to *contingent* propositions, where the entailment relation could still be what makes one such proposition part of another.
16. On balance, then, I think entailment is probably the best candidate for being what determines which pairs of contingent propositions our part–whole concept applies to,

while set-inclusion is probably the best candidate for determining which pairs of properties it applies to.

17. In short, ‘horses for courses’, which is the real moral I want to draw: the relation that makes one entity part of another may vary from one kind of entity to another – a conclusion that Armstrong’s other instances also support
18. What, for example, is the part–whole relation that makes New South Wales part of Australia? We could, I suppose, take that relation to be set inclusion if we identified these geographical entities with sets of points on the earth’s surface.
19. But which sets? Obviously: the sets of all and only those points that lie within the spatial boundaries of New South Wales and Australia respectively. But then the relation that will really determine which pairs of regions of the earth’s surface are related as parts to wholes will not be *set* inclusion but *spatial* inclusion.
20. Similarly with the Terror and the French Revolution, only here we shall need *temporal* as well as spatial inclusion to determine which pairs of historical events are related as parts to wholes.
21. And even these will not be *enough*; since not everything that happened within France during the French Revolution is a *part* of that revolution.
22. For example, the flights of birds across France at that time were not, nor were other natural movements, births and deaths of wild animals. And if some human deaths were parts of the revolution, most were not, and nor were the births of French children.
23. What else it takes to make an event part of the French Revolution, besides spatio-temporal inclusion, is hard to say precisely. But for my purposes I think it will do to say that it has to be what I shall call a *working part* of the revolution.
24. By this I mean that no event is a part of the French Revolution unless some relevant features of the revolution ‘as a whole’ depended on that event *causally*, as they certainly did on the events that were the working parts of The Terror
25. This causal constraint on part–whole relations between events seems to me quite general. Take the more cheerful example of several people having a three-course meal.

- This event – the meal – looks like a whole, with its courses as temporal parts, which in turn have their consumption by the diners as spatial parts.
26. Here too not every event within the meal's spatial and temporal confines is a part of it. Unnoticed movements of insects across the table are not, and nor is every breath the diners take, even if their drinking is.
 27. So here too the extra necessary condition seems to be causal. The only events that *are* parts of the meal are those that make some causal contribution to relevant properties of the meal as a whole.
 28. And as for events, so for so-called *continuants*, or *things*, as I shall call them for short, by which I mean not only inanimate things, both natural and man-made, but plants and animals, including people, and all the other familiar entities that provide the most obvious and least contentious examples of part-whole pairs.
 29. 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.
 30. This is not of course to say that *all* an artefact's parts are things it's assembled from. Artefacts like statues that are made, not by being assembled, but by being carved out of solid blocks of wood or stone, still *have* parts.
 31. 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.
 32. 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.
 33. And as for artefacts, so for natural things like rocks and planets, as well as plants, animals in general, and people in particular. Their parts too we identify with their *working* parts, thereby excluding, for example, radio waves passing through them.

34. All this having been said, let's return to Armstrong's list of apparent part-whole pairs: the proposition p and the conjunction $p\&q$; the property F and the property $F\&G$; the classes 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.
35. Notice again how *heterogeneous* this list is. The entities that it says are related as parts to wholes are pairs, respectively, of propositions, properties, classes, geographical regions, events and things.
36. But equally striking, given this heterogeneity, is the *homogeneity* of each of these pairs. In no case is the whole different in kind from its part. Properties and propositions aren't paired with each other, geographical regions aren't paired with classes, things aren't paired with events, and so on.
37. There are indeed 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 can't be made parts of $\langle Fa \rangle$ just by the entailment relation that makes one proposition part of another.
38. Still, these and other apparent exceptions don't alter the fact that most part-whole relations relate entities of the same kind. And the best way to explain both the rule and the exceptions to it is to allow for a *multiplicity* of part-whole relations ranging from spatiotemporal inclusion plus causal dependence to set inclusion, entailment, and whatever relation gives propositions their non-propositional parts.
39. But of all these examples it's events and things that *really* make the case against postulating a single part-whole relation. For as we've seen, for one event or thing to be a part of another, at least *two* necessary conditions must be met.
40. One is spatiotemporal inclusion, and the other is some kind of causal dependence of the whole on the part. And whatever we take causation to be, these two conditions are manifestly independent, since neither is reducible to the other, or to any of the relations that make one proposition, property or class part of another.

41. Now maybe this shows, as some philosophers think, that we have more than one part–whole *concept*. That seems to me a moot but unimportant point, just as it is, for example, whether the material differences between light, sound and water waves – by which I mean *surface* waves – show that we have more than one concept of a *wave*.
42. But what the existence of different part–whole relations certainly does show is that, as I said 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.
43. For in the wave case, all something has to do to *be* a wave, whatever it's made of, is to satisfy a wave *equation*. It's that *formal* fact about light, sound and water waves that makes us call them all *waves*.
44. This doesn't alter or conflict with the fact that light waves are features of *light* – i.e. of electromagnetic radiation – sound waves are quite different features of *air*, and water waves are equally different features of *water surfaces*.
45. On the contrary. Being features of quite different things is precisely what allows us, if we wish, to apply a single formal wave concept to all these different features, because it's what stops them conflicting. It's what, for example, stops sound waves also being light waves of a different frequency.
46. 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 because they, is the fact that they have different *relata*.
47. This then is why *mereology*, the general study of part–whole relations, is usually – and I think rightly – regarded as a formal rather than a material study. But this way of reading it exposes mereologists to an occupational hazard – that of trying to decide *a priori* what the formal properties of part–whole relations are.
48. For while all part–whole relations must share many if not most of their formal properties -- or they would no more fall under a single concept, or family of concepts, than do river banks and high street banks – it doesn't follow that they share *all* of them.

49. Take for example 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.
50. In fact nearly all systems of mereology assume that part-whole relations *are* transitive. But mereologists can't justify this assumption *a priori*, without checking the transitivity of all the very different relations we wish to call part-whole relations..
51. When we do that, we find that most of them are in fact transitive. Entailment, class inclusion and spatiotemporal inclusion, for example, 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 classes or regions of space, time or space-time.
52. But whether *causal dependence* is transitive is another matter, especially when the dependence is not deterministic. And whether it is or not, that's a question about *causation*.
53. So if, as I've argued, events or things with parts must, among other things, be causally dependent on those parts, then whether transitivity holds here isn't a question we can answer independently of our theory of causation. And then the answer that theory gives will affect what parts we can take events and things to have.
54. If, for example, causal dependence *is* transitive, then 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 isn't transitive, then my parts will only include those of their parts on which I as a whole depend causally.
55. In short, the formal properties of our part-whole concept must be derived from those of our part-whole relations, not the other way round. We can't decide them by consulting *a priori* intuitions about parts and wholes in general.
56. Trying to do that is like trying to decide *a priori* whether waves are longitudinal or transverse, i.e. whether they oscillate in the direction the wave moves in, or at right angles to it: a hopeless project, since different waves oscillate in different ways, sound waves being longitudinal and light and water waves transverse.

57. Admittedly 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, more or less *a priori*, that all waves need a material medium, when the fact is that some do and some don't.
58. Sound waves, for example, do need a material medium: air or water, or some other material fluid; just as water waves need a material boundary, like that between air and a body of water that's being held down by gravity.
59. Light, on the other hand, needs *no* such medium, unless you count electromagnetic fields, which won't do what's wanted, namely provide something for the speed of light to be *relative* to, in the way that the speed of sound and water waves is relative to the air or water they're waves in. Hence the late nineteenth century search for a material *ether*, which we now know doesn't exist.
60. Bearing this cautionary tale in mind, let's return to mereology and, in particular, to the other widely – though not universally – accepted principle of *unrestricted composition*, which says that *all* pluralities of entities compose wholes that have those entities as parts.
61. (Whether these entities are *all* the parts of this whole will of course depend on whether *they* have parts and, if so, whether their parts are also parts of the whole. That, as we've seen, is another question, which for present purposes we can ignore, just as we can ignore the degenerate case of wholes with just one part, namely themselves.)
62. Well, with that caveat, is the principle of unrestricted composition true? The answer, as with transitivity, may – and I shall argue does – vary from one kind of entities to another. For some kinds composition *is* unrestricted and for others it isn't.
63. First, a bit of standard terminology: in mereology any whole is called the *sum* of its parts. That's true by mere definition – it's simply what the term 'sum' means here – and it doesn't entail that composition is unrestricted, since it doesn't entail that *any* entities either *have* or *are* sums other than themselves.
64. Of course, as with transitivity, unrestricted composition *does* hold for several kinds of entities: otherwise the principle would never have been adopted. For example, it

certainly holds for sets and classes. But that's only because it's obvious what we should assume the sum of two or more sets or classes, such as the class of women and the class of men, to be: namely their *union*, in this case the class of which all and only women and men are members.

65. This is the assumption that makes the principle of unrestricted composition true of classes, by making it follow from the fact that any number of classes *has* a union, which is also a class. And however uncontentious this assumption is, it takes more than mereology to justify it, because it's an assumption about *classes*, not just about sums.
66. 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 and any number of propositions *have* a conjunction, which is also a proposition, is what makes unrestricted composition apply to propositions.
67. *Properties* are less obvious, since it's a moot point whether the only credible candidates for being sums of properties, namely *conjunctive* properties like *F&G*, exist. Many - perhaps most - philosophers think they do, often because they identify a property with the class of all its actual or possible instances.
68. Since I don't do this - I think properties are *universals* - I can, and in fact do, deny that there are any such entities as conjunctive properties, for reasons I needn't go into here. And if I'm right, then there is *no* non-trivial composition of properties: properties have no properties other than themselves as parts.
69. I shan't however argue that point here. The point here is that, as with transitivity, whether unrestricted composition applies to properties depends on answers to questions about *properties*, answers that can't be established by a theory of parts and wholes in general.
70. How about regions of *space*, *time* or *spacetime*? Surely, if we identify such regions with the classes of the spatiotemporal points they contain, that will make their composition as unrestricted as that of classes in general is?
71. Not necessarily. Take two disconnected lines, one *here* and one *there*, or two disconnected time intervals, say 2005 and 2007. Do these lines and time intervals have

- sums? If they do, these sums are certainly not connected, and I see no non-question-begging reason to suppose they exist at all.
72. For what, if not the principle of unrestricted composition, makes the existence of any two lines or intervals entail the existence of a third line or interval that is their sum?
 73. To that rhetorical question I see no good answer, except the theoretical utility of assuming unrestricted composition. But then the question is how useful this assumption is, not to a *general* theory of parts and wholes, but to theories of regions of space, time or spacetime.
 74. And for myself, I can't see what assuming unrestricted composition does for such theories except to multiply the entities they postulate beyond both necessity and credibility.
 75. The only regions that I think such theories need to postulate are internally connected ones, which gives them quite enough extended parts and wholes. The loss of unrestricted composition in this case seems to me no loss at all.
 76. Finally, and most importantly, what about *things* and *events*? Is *their* composition unrestricted? Do any two of them, however remote and unconnected, compose a third? Is there really a sum of me and the Eiffel Tower, or of my birth and the Big Bang? Well, on the face of it, obviously not.
 77. Why then do most mereologists *accept* the principle of unrestricted composition of things and events? There is a formal reason, namely that the principle 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.
 78. But models are 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 doesn't show that it's true of actual geodesics, and it isn't.
 79. Similarly with mereology. What we – or at any rate I – want to know is whether the principle of unrestricted composition is true of *actual* things and events, not whether it

assuming it generates a powerful and elegant theoretical model of them. We need a better theoretical reason than that to counter the principle's apparent falsity.

80. One such reason is the utility in semantics of having singular terms for referring to arbitrary collections of entities, terms which words like 'sum', 'set', 'class' – and of course 'collection' – can be used to form.
81. However, the *semantic* utility of singular terms like 'the set of Xs' or 'the sum of Xs' doesn't show that in reality, as well as the Xs that these terms provide useful ways of talking about collectively, there *are* any such entities as the *set* of Xs or the *sum* of Xs.
82. There is, moreover, a crucial semantic difference between sets and sums that makes set theory more autonomous than mereology. The difference is that whereas sets are always of a different semantic kind from whatever they're sets of, sums are not. Sets of events, for example, unlike sums of events, are not events.
83. 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've seen, on what they're sums of.
84. So 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.
85. 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.
86. But this doesn't of course show that *any* two or more events or things *have* a sum, i.e. that unrestricted composition applies to these kinds of entity. So why should we think it does?
87. Why, to take two other examples, should we believe that, as well as the first and second world wars, there is a third event that is the sum of these two events (and perhaps of some or all of their parts)? Or that, besides everyone in this room, there is another thing that is the sum of all of us (and perhaps of some or all of our parts)?

88. But that of course is just another rhetorical question, which gets its rhetorical force from intuitions about what there is and what there isn't, intuitions that defenders of unrestricted composition could dismiss as part of a mere 'folk ontology' that may be superseded by a better theory of parts and wholes.
89. And I agree: mere intuitions about examples cannot defeat good arguments for the theoretical utility of principles like unrestricted composition.
90. But the question is: theoretical utility for what? As I've already argued about other principles, like transitivity, what matters is not whether unrestricted composition is useful to mereology in general, but whether it's useful to a theory of this application of it.
91. And the fact is that it isn't. On the contrary, assuming that any two or more events or things have a sum causes far more theoretical trouble than it saves, as a well-known article by David Lewis, 'Many, but almost one', unintentionally shows.
92. This paper tackles problems posed for mereology by *vagueness* in what parts things have, problems whose analogies for events I think are sufficiently close and obvious that we needn't discuss those separately. From now on, therefore, I shall stick to things.
93. In his paper, Lewis demonstrates the problems that vagueness generates in the case of a *cat*, Tibbles, many of whose loose hairs are borderline candidates for being among its parts.
94. So now consider two sums of its parts, which differ only in that one of them includes one of these hairs and the other does not. Each of these sums is an equally good candidate for *being* Tibbles; and yet they can't both *be* Tibbles, since no hair can both be and not be a part of one and the same cat.
95. Similarly for all the other borderline candidates for being parts of Tibbles, which gives us thousands of sums, all of them equally good – and equally incompatible – candidates for being that cat.
96. What then should we say, given that each of these 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.

97. But we can do better than that, *if* we deny the principle of unrestricted composition, which is what produces Lewis's problem in the first place, by entailing the existence of sums of any number of Tibbles's parts.
98. So if we deny the principle, 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's a vague matter how many hairs someone has to have before they cease to be bald. But this vagueness need only be semantic: it needn't be ontological.
99. Now I don't claim to know what the best solution is to the problems posed by the vagueness of predicates like 'is bald'; I don't. What I do claim is that whatever that solution is, it should apply equally well to the vagueness of predicates like 'is a part of Tibbles'.
100. For whatever the pros and cons of rival views of vagueness, none of them seems to me as absurd as the multiplication of cats and other entities produced by applying the principle of unrestricted composition to things.
101. For that absurdity is worse, far worse, than Lewis admits. For if unrestricted composition applies to things in general, it will apply not only to all parts of *cats* but to all parts of *hairs*, since each hair will also have many borderline candidates for being one of its parts, such as the dead or dying cells on its surface.
102. This will turn each apparently single hair into thousands, thus multiplying a thousand-fold Lewis's already incredibly large number of cats. And then of course every cell of every hair will also have borderline candidates for being parts of it, such as the molecules loosely or temporarily attached to its surface.
103. 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 extending to infinity, thus not even satisfying another basic mereological assumption about things, namely that the parts of things must be *inside* them.

104. In short, not only is unrestricted composition not a useful assumption in a mereological theory of the ontology of things, it may well be fatal, since it makes that theory incompatible with modern microphysics.
105. For while modern microphysics has its conceptual problems, and is certainly neither infallible nor complete, there is overwhelming evidence for the probabilistic features of it that conflict with an unrestricted mereological composition of things.
106. And given this conflict, the problems that vagueness makes it generate, the lack of any other theoretical argument for it that I can see, *and* its apparent falsity, I conclude that the *unrestricted* principle of unrestricted composition is false. The mereological composition of things is *not* unrestricted.