A RELEVANCE CONSTRAINT ON COMPOSITION

David Vander Laan

[This is an Accepted Manuscript of an article published by Taylor & Francis in *Australasian Journal of Philosophy*, March 2010, available online: http://dx.doi.org/10.1080/00048400902941299.]

Whether certain objects compose a whole at a given time does not seem to depend on anything other than the character of those objects and the relations between them. This observation suggest a far-reaching constraint on theories of composition. One version of the constraint has been explicitly adopted by van Inwagen and rules out his own answer to the composition question. The constraint also rules out the other well-known moderate answers that have so far been proposed.

1. Introduction

The question 'Under what conditions do some objects compose another?' has increasingly been recognized as a central question for the ontology of material objects.[1] But as is often the case with philosophical questions, there has emerged no consensus on how best to answer it, or even how best to investigate it. Still, the answers so far proposed are relatively few. There are the extreme answers that objects either always compose something or never do. Between these extremes is Peter van Inwagen's answer: objects compose something if and only if they are 'caught up in a life' [1990]. Joshua Hoffman and Gary Rosencrantz suggest that objects compose something when they are either functionally united or rigidly bonded [1997].[2] And then there are a few conspicuous non-answers. One is the thought that the question has no answer, or at any rate no enlightening general answer, none that can be expressed without, in effect, listing all the composite objects there are [Markosian 1998]. Another is the hope that there is a suitable moderate answer unknown to us [Smith 2006], perhaps one that is consistent with a mereology of common sense [Simons 2006; Sanford 1993].

One of the challenges facing the philosopher who tackles the composition question is finding any plausible general principles on which to base her answer. Most philosophers would prefer an answer consistent with the existence of animals, for example, but it is not obvious what relation between an animal's parts might be so much as relevant. 'What sort of unity might be necessary and sufficient for things to compose something?' Eric Olson has recently asked. 'It is surprisingly hard to say. I don't know of any answer that has much plausibility on the face of it' [2007: 225].

Below I will endorse a relevance constraint on answers to the composition question. The constraint has the virtue of being intuitively strong—what I would call obvious had I never been chastened by philosophers' skill in deposing what might seem evident. In a slightly different form, it has been endorsed by van Inwagen. I will argue that it rules out most of the available answers to the composition question, including van Inwagen's own. In fact, it rules out all the moderate answers to the composition question that have been explicitly proposed. The only proposed answers that are consistent with the constraint are the extreme ones: unrestricted composition (or 'mereological universalism') and compositional nihilism.

2. The Internality of Unity

Let's approach the composition question by way of another: Which sorts of facts should we expect to be relevant to whether composition occurs? This question will be somewhat easier to address if we frame it as a question about synchronic unity.

To say that non-overlapping objects compose something is to say that some individual object is their sum or fusion.[3] (If it is possible for objects to coincide, more than one object may have this feature.) A composite thus has unity in the sense of *oneness*, the property of being one thing, and its parts have unity in the derivative sense of being *united* as parts of the same object. The conditions under which a plurality of objects compose something are the conditions under which an individual composite object has the internal unity necessary and sufficient for its existence.

What sorts of facts, then, might be relevant to whether an object has the requisite internal unity? Internal facts, presumably. We naturally look to the object's intrinsic properties and the relations between its parts. In fact it is difficult to see just how external facts could be relevant. Granted, it would not be surprising if the internal unity of a composite object were caused (at least in part) by the movement of some object or objects disjoint from it. But what we are after is not the efficient cause of an object's internal unity, but rather the conditions under which it obtains, and these would seem to be conditions of the object in question and its parts. Only internal facts are relevant to whether an object has unity.

This thought, that the unity of composites is internal, suggests a more specific version of the general principle.

3. The Constraint

In *Material Beings* Peter van Inwagen adopts a constraint on his view of composition that I will call 'Locality.'

> Locality: The only facts relevant to whether certain objects compose something are facts about their intrinsic properties and the relations that hold between them. In particular, whether objects compose something does not depend on what occurs outside the region they occupy.[4]

We may summarize by saying that synchronic unity conditions must be local.

For illustration, consider Grouping, a hypothetical answer to the composition question that avoids what many would see as a disadvantage of Contact, one of the possible answers that van Inwagen considers and rejects. According to Contact, objects compose something if and only if they are in contact with each other. But given the distances between subatomic particles, its consequent is almost never true. We might imagine a disappointed philosopher looking for a way to loosen the conditions for synchronic unity. A frown, a furrowed brow, and then this:

> Grouping: The *x*s compose something iff the *x*s are pairwise related by the ancestral of being closer to each other than any of the *x*s is to anything disjoint from each of the *x*s.

What Grouping says is that objects compose something when they are grouped together apart from other things. They need not be 'bunched,' each one relatively close to each of the others, provided they are linked in chain-like fashion by a closeness that no link has to anything outside the chain. If outer space were a perfect vacuum, Grouping would imply the existence of an object we could reasonably call 'the solar system.' It would also imply the existence of a larger object comprising the solar system and the postulated Oort cloud, and a smaller object consisting of the Earth and its moon.

Pace our imaginary philosopher, Grouping does not have a lot to recommend it.[5] What it illustrates for us, though, is how an answer to the composition question may run afoul of Locality. Whether the *x*s are closer to each other than to other objects depends on the position of other objects. Sedna is a dwarf planet that orbits the sun well beyond Pluto and the Kuiper belt.

Normally it is far enough away not to be included in the solar system according to our Grouping-based characterization. Let the *x*s be the simples[6] that compose the solar system so construed. There is a point in Sedna's eccentric orbit at which it comes close enough to the objects of the Kuiper belt that the *x*s no longer compose an object according to Grouping since they are not closer to each other than to any other object. If Grouping were true, then whether certain objects composed something would depend on facts about *other* objects, contrary to Locality.

As advertised, Locality highlights certain respects in which the unity of composites is internal. It expresses the idea that the facts bearing on composition are spatially internal and the idea that such facts involve only the purportedly composing objects and not others. I will assume that the constraint is necessarily true if true.

4. Life and Locality

Locality rules out a variety of the proposed answers to the composition question. For example, it rules out conventionalist answers, i.e., those according to which whether objects compose something depends on which conventions some community adopts. Composite-making conventions would not (in general) be made within the boundaries of a putative composite object.

More striking is the fact that Locality is inconsistent with Life, van Inwagen's preferred answer to the composition question. To see the snag, consider a child who falls and scrapes her elbow. According to Life, certain simples that compose an organism at a moment just after the fall (the simples of some set B) did not compose anything before the fall. What changes is the relationship between the simples of set B and the simples lost with the scrape. That is to say, whether the simples of set B compose anything depends on what occurs outside the region they occupy. In particular it depends on their relations with certain simples that are not among them. Had the simples removed in the scrape not been lost, the simples that in fact composed an organism afterward would not have done so. Thus Life violates Locality. Given Life, whether objects collectively have synchronic unity—whether they compose one thing—is not a strictly internal matter but depends on external facts.

One objection to this line of argument is that the relationship between the simples of set B and the others is not all that changes when the child falls. Realistically, the relations between the simples of set B change as well. But if in any possible world such a loss occurs without change in the internal relations between the members of B, whether by a small miracle or by physical laws slightly different from the actual ones, then Life cannot be true in that world. Grant this modest assumption: if Life is true, it would be true even if a small miracle were to occur or if locally anomalous laws held. But then if Locality is true, Life is not.

This conclusion could be avoided were the defender of Life to adopt a sufficiently broad conception of life, one which included 'non-maximal' lives. He would then be free to claim that the simples of set B composed something even before the scrape and that whether the simples of set B composed something was an internal matter. Unfortunately, this stance would rob him of the most promising solutions to the paradox of undetached parts.[7]

In van Inwagen's well-known example, we are to imagine that Descartes once lost his left leg. 'D' names Descartes (here assumed to be an organism), and 'D-minus' names all of Descartes but his left leg. It looks as if what remains after the accident is D-minus; this proper part of Descartes presumably survives its detachment from Descartes' left leg. But what remains also seems to be D, Descartes himself. D and D-minus are clearly not identical to each other before the accident. They differ in shape and in mass, for example. But then they are not identical to each other after the accident either. So on the face of it, D is identical to D-minus, and, on the face of it, D is not identical to D-minus.

Life avoids the paradox by denying the existence of D-minus. The relevant simples, those we might be tempted to think compose all of Descartes but his left

leg, are not simples whose activity constitutes a life. Thus according to Life they do not compose an object. There is no such thing as D-minus, and the apparent contradiction is avoided. But if the proponent of Life asserts that D-minus *does* exist (and in our unfortunate youngster's case, that the simples of set B compose something before the fall) contradiction again looms.

Other responses to the paradox of undetached parts have been proposed, but not all of them will be palatable. The proponent of Life can hardly be expected to embrace the view that persons are immaterial objects, or tiny material objects lodged in the brain. One of the assumptions motivating Life to begin with is that human beings are biological organisms [van Inwagen 1990: 75-76]. A temporal parts solution might be consistent with Life, though the two would make an anomalous pair.[8]

Another possible response posits co-location: first, the co-location of our heroine with a sum of the larger set of simples (let's call that set A), and later, the co-location of our heroine with a sum of the smaller set of simples (B). Co-location itself is no violation of Locality. Locality says that whether objects compose something depends only on local facts, but it does not say how many things the objects might compose, or which. However, relevance considerations of the sort that make Locality plausible also suggest a variation on it:

> Locality*: Only local facts are relevant to whether given objects come to compose something they did not compose earlier (or cease to compose something they did compose earlier).

According to the co-locationist, the simples of set B do not compose the girl before the fall and do compose her afterward. But again the local facts—facts about the arrangement and intrinsic properties of the simples of set B—have not changed in any important way. The co-locationist's response to the paradox of undetached parts fails if Locality* is true.[9]

In my estimation Locality* has less intuitive éclat than the original version, though only a bit less.[10] Whether they already compose something or not, objects that come to compose something acquire a new unity as objects that together make up some particular composite. Having such a unity would not seem to depend on what happens to other, non-overlapping objects.[11]

This is not the place to give a full assessment of the possible responses to the paradox of undetached parts, but we have seen enough to know that a solution will come to most proponents of Life only at a prohibitive cost.

The situation for Hoffman and Rosenkrantz's disjunctive answer to the composition question is similar. There is reason to think that the simples of set B are not a functional unit before the girl's fall. But if functional unity is conceived broadly enough that they do, the paradox of undetached parts looms, and each of the available responses would make a strange bedfellow with the disjunctive answer.

5. Fuzzy Life

Though van Inwagen does not address the foregoing objection directly, he does consider similar "missing simples" cases when giving his reply to the problem of the many, which can be posed as a difficulty stemming from Locality. If there is a solution to the problem that Locality poses for Life, we might hope to find it here. Let's sketch out the problem of the many and van Inwagen's proposed solution.

The problem is that there is evidently an abundance of overlapping collections of simples whose members' activities constitute lives. Take any collection of simples whose members do constitute a life, and suppose that one of them were somehow removed without affecting the activity of the others. That activity would then constitute a life, and, according to Life, the simples involved in it would compose an organism. Now we apply Locality. Whether given simples compose something depends only on local facts, so those simples that would

compose something were some other nearby simple taken away do in fact compose something. There are many such collections of simples, to say the least. So Life and Locality imply that there is a proliferation of organisms where we would normally be inclined to identify only one.

The solution van Inwagen offers is a modification of Life in terms of fuzzy sets, i.e., sets that allow degrees of membership other than definite membership and definite nonmembership. The modification can be expressed as follows.

> Life*: The simples belonging to a fuzzy set F compose something iff their collective activity constitutes a life, i.e., each of the simples is a member of F to the degree that it is "caught up in" that life.

Each simple, van Inwagen says, is caught up in a given life to a degree. Of the various fuzzy sets of simples whose members purportedly compose overlapping organisms, only one has this feature: each of its members is caught up in a certain life to exactly the degree that it is a member of the fuzzy set. And so only one of the competing collections of simples in fact composes anything.

For the sake of argument, grant that it makes sense to speak of simples being caught up in a life to a degree and that fuzzy sets are an appropriate mathematical tool for representing those degrees. Life* is nonetheless inconsistent with Locality in just the way Life is.

Again, a girl falls and sustains a slight abrasion to her elbow. According to Life*, the simples of a certain fuzzy set F compose her immediately after the fall but did not compose her beforehand. As far as their internal properties and relations are concerned, however, the simples of F were then suited to compose an object. Locality tells us that facts about the internal properties and relations of the simples of F are the only ones relevant to whether they compose something, so before the injury they did compose something, contrary to Life*. Though van Inwagen does offer reasons for believing that Life* can avoid the problem of the many, none of these suggest that Life* can be reconciled with Locality.

Faced with the inconsistency between Locality and Life (in both the original and fuzzy formulations), the defender of the latter might perhaps choose to dispense with the former and carry on without it. Perhaps the cost of doing so will be seen as quite manageable. After all, many of the views Locality rules out (such as conventionalist theories of composition) could be rejected on other grounds.

As I see it, the primary difficulty with this is simply that Locality is intuitively strong. Here is how van Inwagen introduced his tenth constraint:

> Whether certain objects add up to or compose some larger object does not depend on anything besides the spatial and causal relations they bear to one another. If, for example, someone wants to know whether the bricks in a certain brickyard make up a composite object, he need not attend to anything outside the brickyard, for no information gathered from that quarter could possibly be relevant to his question [1990: 12].

This is the final word of van Inwagen's decalogue, the last of the theses he says he would have accepted even had he been unable to think of arguments for them. It is fair to say that Locality, even in his somewhat stronger version, has an appreciable intuitive force for van Inwagen[12]—and, I suspect, for many others as well. Some philosophers, drawn to Life as an answer to the composition question or suspicious of the alternatives, will consider rejecting Locality. If the arguments of this paper are successful, they will have to. However, the appeal that Locality apparently has gives rejecting it a significant cost, one that cannot be off-handedly dismissed.

6. Dynamic Systems Answers

Toward the goal of finding out to what extent can the above arguments be generalized, consider two hypothetical answers to the composition question, each of which, like Life, accommodates the existence of organisms.

> Replication: The *x*s compose an object iff, in virtue of their dynamic relationships, the *x*s cause (or have the capacity to cause) other objects (possibly including some of the *x*s) to take on a form of the same type.
>
> Homeostasis: The *x*s compose an object iff their collective activity is in dynamic equilibrium.[13]

Replication implies the existence of entities not countenanced by Life, such as DNA, viruses, and perhaps even memes (or their physical cognates). Homeostasis implies the existence of hurricanes, ecosystems, and galaxies. Collectively, let's call these 'dynamic systems' answers.

How do the dynamic systems answers fare vis-à-vis the charge that they violate an intuitive relevance constraint on composition? In an experimental spirit, let's cast the argument as we did before.

> The girl exists, assuming only that she is an organism. And assuming that organisms are long-lived, we may add that she exists both before and after her fall. The simples of set B compose her just after the fall, but they do not compose anything before the fall. Beforehand they are only some of the simples whose activity is in dynamic equilibrium, or some of the simples that are together capable of passing along their form. So before the fall they compose nothing, and after the fall, without any significant change in their internal relations, they do compose something, contrary to Locality.

This argument does not fare quite so well as the original. Its most dubious premise is that the simples of set B did not compose something beforehand. Suppose first that Homeostasis is true. Didn't those simples have a strong tendency to maintain a stable, dynamic activity before the fall? Some neighboring simples were causally involved, but the same is true in nearly every case of dynamic equilibrium.

Similarly, in the case of Replication the post-fall simples arguably had all along the capacity to cause other simples to assume a form of the same type. (Let's call this the capacity for replication, taking care to treat it as a capacity of the many simples regardless of whether they compose a whole.) We have not made any attempt to say how broad the relevant type is. If it is broad enough that, before the fall, both the simples of set A and the simples of set B instantiated the same type, then the latter did then have the capacity for replication. Those very simples, in nearly the same arrangement, had the capacity for replication after the fall, clearly enough, and we can hardly suppose that the fall conferred that capacity on them. So before the fall they composed something. If, on the other hand, the simples of the two sets instantiated different types before the fall (*organism* and *near-organism*, say), then the simples of set B then had the capacity for replication because they had the capacity to cause other simples to assume a near-organismal form. In either case, the simples of set B composed something all along, and the doubtful premise of the loss-of-parts argument is false.

So the dynamic systems answers can, with some plausibility, escape the charge that they are inconsistent with Locality. But in each case, the answer's conjunction with Locality comes at the cost of embracing a multitude of overlapping objects. If the notions of replication and equilibrium are broad

enough to include the simples of set B before the fall, then they are broad enough to include many, many collections of simples. That's a problem because it uncomfortably restricts the range of viable responses to the paradox of undetached parts. But it is also a problem if we are to identify ourselves with a sum of one of the collections. Of the sums in my vicinity, an enormous number of them are equally good candidates for being me; it doesn't look as if there are any criteria that might select one of the many for the post. Even if one of them could somehow be identified as me, there remains the multitude of other sums, each apparently thinking thoughts indistinguishable from my own, each a person as fully as I myself am. In my judgment, the cure is worse than the disease.

Returning to the question of generalization, the dynamic systems answers considered here are illustrative; it seems likely that any answer to the composition question that identifies composites with the dynamic systems of a class that includes organisms will encounter similar obstacles. Assuming that the composition question has a general answer, Locality and whatever strong reasons for believing in organisms there may be[14] together generate a paradox. If there are no strong reasons, Locality tips the scales against mereologically dynamic organisms.

7. Naïve Mereology

What we have seen is that Locality rules out the best-known moderate answers to the composition question and their closest cousins as well. The inconsistency can be avoided only by adopting extreme measures: (1) an awkward solution to the paradox of undetached parts and (2) the existence of a multitude overlapping each person.

Once the most promising moderate answers have been crossed out, what options remain? And what implications does the constraint have for the partial answers and non-answers found in the literature? Very briefly, let's consider naïve mereology.

Naïve mereology accepts the objects and parts recognized by common sense. On this view our ordinary ways of thinking and speaking about objects are, if not authoritative, at least fairly reliable. The naïve mereologist does not propose necessary and sufficient conditions for composition—the ontology of common sense is too heterogeneous for that—but she may discern some regularities in the structure of the world. Take for example the view of Peter Simons, who says that our theorizing about wholes and parts should be 'in accord with the real world.' It is no mere coincidence that organisms and artifacts are among the composites we find in the 'real world,' since, Simons seems to hint, having a function is sufficient for being a real object [2006: 597, 609].

A naïve mereology that includes a general principle of this sort will be vulnerable to the objections raised above, since it will violate Locality in the way that Life does. The problem is even more acute in this case since the strategies of embracing the problem of the many and finding a solution to the paradox of undetached parts (other than van Inwagen's) would require significant departures from common sense.

There is another path for the naïve mereologist. Rather than accepting a general principle that says that functioning (or life or replication or equilibrium or some other feature) is sufficient for composition, she might say that the organisms, artifacts, and other objects of everyday life just happen to exist. That is, she might accept brutal composition, the thesis that there is no general answer to the question of when objects compose something.[15] Whatever the merits of this thesis might be, though, its conjunction with naïve mereology invites a pointed question: Why believe in naïve mereology? If common sense included some insight into a shared feature of ordinary objects that might plausibly be thought to ground composition, the pointed question would have an answer. But if there is no general answer to the composition question, there is no shared feature of composites that might lead us to naïve mereology. Why then would we think that the composite objects are the ones that we happen to recognize in our

pragmatic, everyday thinking? Brutal composition better supports mereological agnosticism than it does naïve mereology (or any other detailed ontology).

Those who find themselves with the conviction that unity or oneness is an internal feature of whatever has it thus have strong reasons to reject nearly all moderate answers to the composition question. Barring some hitherto unknown answer that satisfies Locality, the only option that remains is mereological extremism.[16]

Westmont College

REFERENCES

Burke, Michael 1994. Dion and Theon: An Essentialist Solution to an Ancient Puzzle, *Journal of Philosophy*, 91/3: 129-139.

Burke, Michael 1996. Tibbles the Cat: A Modern Sophisma, *Philosophical Studies* 84/1: 63-74

Hoffman, Joshua and Rosencrantz, Gary 1997. *Substance: Its Nature and Existence*, London: Routledge.

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

Markosian, Ned 1998. Brutal Composition, *Philosophical Studies* 92/3: 211-249.

Olson, Eric 2007. *What Are We? A Study in Personal Ontology*, Oxford: Oxford University Press.

Rea, Michael, ed, 1997. *Material Constitution: A Reader*, Lantham, MD: Rowman and Littlefield.

Rea, Michael 2000. Constitution and Kind Membership, *Philosophical Studies* 97/2: 169-193.

Rosen, Gideon and Dorr, Cian 2002. Composition as a Fiction, in *The Blackwell Guide to Metaphysics*, ed. Richard Gale, Oxford: Blackwell: 151-174.

Sanford, David 1993. The Problem of the Many, Many Composition Questions, and Naïve Mereology, *Nous* 27/2: 219-228.

Sider, Theodore 2001. *Four-Dimensionalism,* Oxford: Clarendon.

Simons, Peter 1987. *Parts: A Study in Ontology*, Oxford: Oxford University Press.

Simons, Peter 2006. Real Wholes, Real Parts: Mereology without Algebra, *The Journal of Philosophy* 103/12: 597-613.

Smith, Donald 2006. The Vagueness Argument for Mereological Universalism, *Pacific Philosophical Quarterly* 87/3: 357-368.

Thomasson, Amie 2007. *Ordinary Objects*, Oxford: Oxford University Press.

van Inwagen, Peter 1981. The Doctrine of Arbitrary Undetached Parts, *Pacific Philosophical Quarterly* 62/2: 123-137.

van Inwagen, Peter 1990. *Material Beings*, Cornell: Cornell University Press.

[1] The question is van Inwagen's Special Composition Question. Since I will not need to contrast it with his General Composition Question, I will simply call it 'the composition question.'

[2] The suggestion is not entirely explicit. Some other authors (Peter Simons, for example) suggest that functional unity is at least a sufficient condition for composition [2006].

[3] For the standard definitions of 'overlap,' 'compose,' and 'sum,' see [Simons 1987].

[4] van Inwagen's own statement is slightly stronger: 'Whether certain objects add up to or compose some larger object does not depend on anything besides the spatial and causal relations between them' [1990]. He goes on to endorse a similar constraint on diachronic identity as well. I will not be concerned with the diachronic version.

[5] Among other things, we might note that Grouping contains the mereological term 'disjoint' on the right side of the biconditional, though not in such a way as to make the answer vacuous.

[6] For convenience I will adopt van Inwagen's assumption that every composite is a sum of simples, i.e., objects without proper parts.

[7] Eric Olson has recently called it 'the Paradox of Increase' [2007]. See especially van Inwagen's seminal discussion [1981].

[8] As a class, temporal parts advocates tend to be mereological universalists. For two salient examples, see [Lewis 1984: 212-213] and [Sider 2001: 121-132].

[9] Notice that this constraint also rules out those theories according to which the constituents of an object can come to compose a single, different object without any change in their intrinsic properties and internal relations. For examples of such theories see [Burke 1994], [Burke 1996], and [Rea 2000].

[10] In an earlier draft of this paper I offered a constraint we may call Locality**: Only local facts are relevant to whether certain objects compose any *particular* object. Applying this principle to the current example, it cannot be that certain objects compose the girl after her injury if they had the same intrinsic properties and internal relations before the injury and did not compose her then.

An anonymous referee worried that Locality** would be susceptible to a counterexample of this sort. Suppose there are two possible worlds, w_1 and w_2, with exactly the same simples in exactly the same distribution at some time t. In w_1 Romulus is composed of simples x_1-x_n at t and his brother is composed of simples y_1-y_n at t. In w_2 these roles are reversed: Romulus is composed of y_1-y_n and his brother of x_1-x_n. This is possible, on this line of thinking, because the simples have different histories in the two worlds. In w_1, at some time well before t, Romulus began the process of incorporating the simples x_1-x_n (by a strategic diet plan, say) while his brother began to take in y_1-y_n. In w_2 Romulus began incorporating y_1-y_n instead, until at t he was composed of precisely those simples and, as it happened, all the simples in question enjoyed just the distribution they had in w_1 at t. Thus in one world the simples x_1-x_n compose Romulus and in another they do not, though their relations with each other and their intrinsic properties (we may assume) are the same in both cases. It follows that Locality** is false.

In the text I opted to circumvent the hurdle by using a different constraint (happily, I think the referee's worry prompted me to make an improvement), but not because the proposed scenario is clearly possible. I will argue in the final sections that Locality is difficult to reconcile with the existence of mereologically dynamic organisms, and the nonexistence of such organisms would undermine counterexamples of the kind we have here. Those who have their doubts about the putative possibility it describes may find that they prefer Locality** to its replacement.

[11] Of course there are a variety of other criticisms of co-locationism to which we might appeal if necessary [Olson 2007: 60-75].

[12] In personal communication van Inwagen reports that he remains inclined to hold the constraint despite the sort of argument I have offered. He doubts that it is possible for the simples composing a recently-diminished organism to be properly among those that composed it earlier and intrinsically related as they were then; but his doubt does not reach the status of a settled opinion on the matter.

[13] 'Dynamic equilibrium' might be defined in relatively broad or narrow ways, allowing for different degrees of instability. A precise characterization is unnecessary for present purposes.

[14] Though this is not the place for a thorough assessment of the arguments for the existence of organisms, it is worth noting that the question of what *we* are is controversial [Olson 2007], and so an appeal to our own existence is not sufficient on its own. Nor is a simple appeal to common sense. (But see, e.g., the arguments for commonsense conclusions given by Amie Thomasson [2007].) Nor is it clear that biology gives us reason to believe in composite, individual organisms as opposed to many-object systems with a relatively high degree of functional closure [Rosen and Dorr 2002].

[15] Note that brutal composition is not an answer to the question 'Under what conditions do some objects compose another?' but the denial that a general answer exists. *A fortiori* it is not a moderate answer. In principle one could adopt both Locality and brutal composition, but this would commit one to leaving the composition question unanswered.

[16] Thanks to Peter van Inwagen for clarification of his leanings and to Tom Crisp and the Santa Barbarians for helpful feedback on an earlier draft.