Vagueness, Multiplicity and Parts

This is a preprint of an article accepted for publication in *Noûs* ©2004 Blackwell Publishing

When do objects form a whole? Pre-theoretically we are inclined to think that sometimes objects are parts of a larger object – pages are parts of a book, fingers are parts of a hand, and also parts of people. Other sets of objects we are somewhat inclined to think do *not* form a whole – your fingers and my palm do not make up an object in the way that my fingers and my palm do. Some dust in the Andromeda galaxy and some rock from the centre of Mt Everest do not present themselves as part of a whole in the way that, for example, the cells of an individual redwood tree do.

So much for how things seem on first blush. We can try to improve our grasp of which things make up objects and which do not in various ways – we can assemble a range of judgements about what is part of what and see if there are interesting underlying generalisations, we can see how questions of what is part of what fit into other questions we are concerned about (what a person or an artifact *is*, for example, or theories of creation and destruction). We can take into account unexpected empirical discoveries (when we discovered the atomic nature of objects like trees or rocks or ourselves, with atoms being separated from each other by many atom's diameters, we became less inclined to think objects must be in contact in order for them to make up a whole). We can think about how to apply the concepts of part and whole outside the main areas they were initially applied – we can see what happens when we try to apply them to space, or time, or sets, or numbers, or meanings, or events. And we can take account theoretical virtues such as simplicity and unificatory power, preferring simple powerful theories to gerrymandered, *ad-hoc*-seeming generalisations or to mere collections of data points. There may be other channels of investigation as well. In short, we can do metaphysics.[1]

[1] For a recent argument that this sort of activity is misguided, see van Fraassen 2002, chapter 1. I disagree with van Fraassen that the philosophy of mereology has been little more than solving problems by postulation (p 10), or "idle word play" (p 27). The death of the Death of Metaphysics appears to be a protracted one.

One popular metaphysical conclusion about the part-whole relation is that it obeys a principle of *Unrestricted Composition*: for any group of objects whatsoever, there is a whole that they make up, and in a sense *only* they make up – it is a *fusion* of those objects, in a sense shortly to be explained. Unrestricted Composition seems to disagree with commonsense (and it certainly goes well beyond it) – while it allows that there is a whole object whenever commonsense says there is, it says that there are wholes where commonsense does not (there is an object which is my left ear plus the Alpha Centauri system, and it does not include intervening objects in the intervening space, or elsewhere). I think the fact that it (uncontroversially) goes beyond common opinion and (almost uncontroversially) disagrees with common opinion is something that needs to be taken into account in evaluating it, and indeed that it counts against the view somewhat. The argument, as I see it, is whether unrestricted composition has enough welcome features (e.g. its simplicity and scope) to justify us in departing from common opinion in this respect (or even to justify us adopting it as opposed to suspending judgement).[2]

There is an argument, however, that common opinion should be neglected when it comes to deciding whether or not unrestricted composition is true. With our "intuitions", or commonsense judgements, out of the way, the debate shifts dramatically. The argument was originally offered by David Lewis on behalf or unrestricted composition (Lewis 1986 p 212), but opponents of unrestricted composition have also drawn lessons from Lewis's argument (e.g. van Inwagen 1990 pp 213-283). Despite the interest this argument has generated, I think that even granting Lewis's premises, the conclusion does not follow. On the other hand, I think that if the conclusion *did* follow, we would have an argument on our hands with much wider applicability than people have realised – this style of argument would be widely available to block appeals to intuition and common opinion in many areas of fundamental metaphysics. In this paper I will outline Lewis's argument and some of the effect it has had on the debate about Unrestricted Composition. Then I

[2] I remain uncommitted in this paper on the question of how much weight we should give to the part of common opinion that conflicts with unrestricted composition. Much of what we believe on the basis of investigation does not accord with the verdict of common opinion at the beginning of the investigation – it remains to be seen whether our mature theory of parts and wholes will accord with common opinion any more than our mature theories of mechanics or biology accord with our pre-theoretic conceptions of the physical or biological worlds.

will explain why I think that if the argument worked it would prove to be powerful across a wide range of areas, to the point where it should seem too good to be true. I will then discuss Sider's recent and related argument for temporal parts in the philosophy of time (Sider 1997, Sider 2001): Sider also presents a careful version of Lewis's original argument that I will discuss. I will then discuss the main place where I think both Lewis and Sider go wrong, and discuss what follows when we notice this gap in the argument. I conclude with a discussion of some concerns that might be suggested by the arguments of Lewis and Sider, even if the arguments as presented should be resisted.

For what it is worth, I am a fan of unrestricted composition, and a fan of temporal parts as well. So it is not the conclusions per se of Lewis's and Sider's arguments that I am opposed to – it is just that I do not think you can get *there* by this route. I am also prepared to grant, for the sake of the argument, the main principles about the nature of vagueness itself that Lewis and Sider rely on, since even if these principles are correct, their conclusions do not follow.

1. Lewis's Argument For Unrestricted Composition

Lewis admits that our intuitions (or at least some of them) seem to support some sort of restricted composition. However, Lewis thinks there is a reason we should discount these intuitive judgements:

> The trouble with restricted composition is as follows. It is a vague matter whether a given class satisfies our intuitive *desiderata* for composition. Each *desideratum* taken by itself is vague, and we get still more vagueness by trading them off against each other. To restrict composition in accordance with our intuitions would require a vague restriction. It's not on to say that somewhere we get just enough contrast with the surroundings, just enough cohesion... to cross a threshold and permit composition to take place, though if the candidate class had been just a little worse it would have remained sumless. But if composition obeys a vague restriction, then it must sometimes be a vague matter whether composition takes place or not. And that is impossible. (Lewis 1986 p 212)

To reconstruct, the argument is something like this:

(1) The intuitive criteria for restricting composition are vague, but

(2) When composition occurs cannot be a vague matter, so

(3) When composition occurs cannot fit the intuitive criteria.

The first premise of this reconstructed argument seems true. As Lewis says, "[w]e are happy enough with mereological sums of things that contrast with their surroundings more than they do with one another; and that are adjacent, stick together, and act jointly. We are more reluctant to affirm the existence of mereological sums of things that are disparate and scattered and go their separate ways." (Lewis 1986 p 211). Being close enough to each other, being stuck together enough, being homogenous enough, all seem to be a matter of parts fitting vague criteria.

The second premise is much less obvious or uncontroversial, and it is this premise that Lewis supports with further argument:

> The only intelligible account of vagueness locates it in our thought and language. The reason it's vague where the outback begins is not that there's this thing, the outback, with imprecise borders; rather there are many things, with different borders, and nobody has been fool enough to try to enforce a choice of one of them as the official referent of the word 'outback'. Vagueness is semantic indecision. But not all of language is vague. The truth-functional connectives aren't, for instance. Nor are the words for identity and difference, and for the partial identity of overlap. Nor are the idioms of quantification, so long as they are unrestricted. How could any of these be vague? What would be the alternatives between which we haven't chosen?
>
> The question whether composition takes place in a given case, whether a given class does or does not have a mereological sum, can be stated in a part of language where nothing is vague. Therefore it cannot have a vague answer. There is such a thing as a sum, or there isn't. It cannot be said that, because the *desiderata* for composition are satisfied to a borderline degree, there sort of is and sort of isn't. What is this thing such that it sort of is so, and sort of isn't, that there is any such thing? No restriction on composition can be vague. But unless it is vague, it cannot fit the intuitive desiderata. So no restriction on composition can serve the intuitions that motivate it. So restriction would be gratuitous. Composition is unrestricted. (Lewis 1986 p 212)

We can reconstruct a sub-argument for something like premise 2:

a. Vagueness is only in thought and language.

b. Whether composition takes place in a given case can be stated in a part of language where nothing is vague (the language of unrestricted quantification, identity and difference, overlap, etc.)
so
c. There is nothing vague about whether composition takes place in a case.

The conclusion (c) is not quite the same as (2) yet, since (c) says that composition is not vague, and (2) says that composition cannot be vague. We might bridge this gap by speaking of all possible worlds at once in (a) and (b), or we could just supply necessitated versions of (a) and (b). Since (a) and (b) are presumably necessary truths if true at all (at least necessary given that there is vague thought and language), we can help ourselves to a necessary version of (c) – which gives us the conclusion not just that there is no vague composition, but that there cannot be anything vague about when composition takes place. (Of course we may still be able to describe cases of composition vaguely – e.g. "the middle part of that finger is part of a bald man". We can get vagueness in a sentence about composition by employing vague referring terms or vague predicates, as always. The point is, if Lewis is right that this vagueness does not stem from anything about composition, but only from other sources.)

There are several conceptions of vagueness which take it to be a matter of semantic indeterminacy over a range of cases, which are in themselves in principle describable without vague language (for instance as in the relation Lewis envisions between the word "outback" and the relevant precisely demarcated regions of land). The most well-known of these is the approach known as "supervaluationism". The technique of supervaluationism was first developed by van Fraassen 1966 and first applied to vagueness by Fine 1975. According to the standard version of supervaluationism, a vague predicate like "heap" is associated with many different precise extensions – these extensions are called *precisifications*. (You could think of them as possible acceptable sharpenings of the predicate, though this is not mandatory). Supervaluationism relies on a notion of "truth according to a precisification" or "truth under a precisification" – in the case of heap, something is a heap according to a precisification if it falls under the extension. Paradigm cases of heaps will be heaps according to all admissible precisifications, and paradigm non-heaps will not be heaps according to any admissible

precisification. A vague sentence is true if it is true under all admissible precisifications ("supertrue" as it is sometimes called), and the sentence is false if it is false under all such precisifications ("superfalse"). It is neither true nor false otherwise – thus for the orthodox supervaluationist vagueness gives rise to "truth value gaps". A recent influential presentation of orthodox supervaluationism about vagueness and some of its problems can be found in Williamson 1994.

There are, however, other approaches which treat vagueness as entirely in thought and language, and a matter of a relation between a single piece of language and many pieces of the world, where each of those pieces of the world are in principle able to be picked out in a non-vague way (e.g. objects with precise boundaries, predicates with determinate extensions). We may follow Sider 2001 (and, close enough, van Inwagen 1990) in calling the members of this family versions of "the linguistic theory of vagueness". One alternative version is the approach of Theodore Sider and David Braun (Sider and Braun, unpublished), according to which a vague sentence, even one taken as being uttered in a specific context, does not express a single proposition at all, but instead is associated with a multiplicity of precise propositions. A vague sentence is semantically appropriate to say if the "legitimate disambiguations" of the sentence are true, but strictly speaking is neither true nor false. A similar position is advocated by Lewis 1979, drawing on Lewis 1970 – vague sentences are strictly speaking true or false according to a "delineation", and for it to be appropriate to utter a sentence it must be true according to "a large enough part of the range of its delineations", where how large this must be is a matter of context.[3]

Another is the "sub-valuational" response recently defended by Dominic Hyde (Hyde 1997). This approach is a close relative of supervaluationism. Where supervaluationism says that the truth of a vague claim is a matter of truth on all admissible precisifications, subvaluationism holds that a vague sentence is true provided that it is true on *some* precisification. Thus a borderline case of baldness is a case of both baldness and non-

[3] It is sometimes claimed that Lewis defends a supervaluational account of vagueness in either Lewis 1979 or Lewis 1970. Both Williamson 1994 p 146 and endnote 12, and Hyde 1997 p 650 footnote 11 say that Lewis's 1970 view is a supervaluationist one – but it looks like Lewis 1979 did not think so, and it is hard to read what Lewis 1970 says as explicitly commiting his view to supervaluationism, rather than a bivalent "truth relative to a delineation" account.

baldness. This subvaluational story has some drawbacks, which it wears on its sleeve – it's inconsistent about vagueness, and it does not allow that conjunction introduction is valid, but it clearly falls in this family of "linguistic" approaches to vagueness.

Yet another kind of semantic approach that requires multiple precise candidates is exemplified by "contextualist" approaches, two recent examples of which are Soames 1999 and Graff 2000. According to these approaches, there are a variety of lines that could be drawn by different contexts between e.g. the bald and the not-bald, and different lines will be appropriate in different contexts. (Vagueness is to be distinguished from other forms of context-dependence, which may also be present – presumably in terms of *how* vague terms interact with contexts). Here, too, there need to be many precise boundaries associated with a vague expression (boundaries that may be selected by different possible contexts), and here too there need not be any "vagueness in the world" – there need not be objects and properties with vague boundaries and vague applications, as well as all of the precise objects and properties that different contexts will associate with vague terms.

Regardless which "linguistic theory of vagueness" we adopt, the core is that there are many precise meanings associated with a vague expression, and the semantic value of the vague expression depends on how those precise meanings operate. Let us say that a vague sentence is *determinate* if it receives the same truth value according to each precisification (/delineation/possible demarcation of context). For the orthodox supervaluationist, only determinate sentences receive truth values – indeterminate sentences, true on some precisifications but false on others, are truth-valueless. Other theories of vagueness will treat indeterminate sentences differently. Despite noting this variety, and despite the fact that neither Lewis nor Sider are necessarily supervaluationists about vagueness, it will do no harm when considering these arguments to keep orthodox supervaluationism in mind as a candidate account of vagueness.

Much attention could of course be directed at the principle about vagueness that Lewis relies on. For one thing, I have not said very much about Lewis's claim that vagueness in

composition would make a difference to the truth-value of statements "in a part of language where nothing is vague". For now, though, I want to indicate why we should be suspicious that this argument could work by considering what would happen if it was generalised.

2. Too Good to be True? Some Reasons to be Suspicious

Suppose we agree, for the time being, that vagueness is entirely in language: facts about what the world is really like are not vague (and are capable of being captured by a non-vague language, at least in principle, let us suppose). Nevertheless, our intuitive principles of composition draw vague boundaries. So the metaphysical lines drawn in the world cannot be the vague lines drawn by our intuition. So the boundaries of which objects or stages make up mereological wholes have to be drawn elsewhere. This seems to me to be the general thrust of Lewis's argument.

This sort of argument can be generalised. Whenever we are engaged in constructing a picture of how the world fundamentally is, in principle we should not let any of the fundamental metaphysical distinctions be vague – sometimes the best we may be able to do with our representative resources is provide a vague characterisation of the world, but we should think that the "joints in nature" themselves are not fuzzy or indeterminate. (After all, the fuzziness and indeterminacy are in our representation not the thing represented according to this view of vagueness). However, our pre-theoretic intuitions and everyday judgements about the world are almost invariably couched in vague language, and the distinctions they draw will not be the same distinctions as are drawn by our ideal accurate description of the world. Any fundamental metaphysical distinction therefore cannot "serve the intuitions that motivate it", and since our ordinary thought and talk cannot therefore be getting the boundaries right, we are entitled to disregard our vague intuitions about where the metaphysical boundaries are.

This generalisation seems to prove too much. It may be that there is no vagueness in the world, but nevertheless construction of a non-vague but adequate description of the world

can be constrained by vague theory and "intuitions". We believe, for example, that some people are bald. Some possible completely precise distributions of molecules make it true that some people are bald, and others do not. If a metaphysics proposes a complete precise arrangement of the world's features that does not make it true that some people are bald , we have reason to be suspicious of it. Or consider another case. Suppose a Laplacian demon understands complicated, completely non-vague descriptions of the positions of molecules and their directions (suppose the world is a simple corpuscular Newtonian one, if the idea of understanding precise descriptions of precise quantum what-evers is giving you a mental block). The Laplacian demon is presented with a stack of such precise descriptions, and is commanded by her Infernal Overlord "Find me a description of my favourite region!". Since the demon has no idea of her Overlord's tastes in regions she is unsurprisingly nonplussed. But since her Overlord will reward the correct answer, and punish the wrong answer (or no answer) with hellfire, the demon decides she must pick one of the regions described and hope she has it right.

While the demon is reading through the descriptions, she is told by a human, and so irredeemably vague, informant that the Overlord's favourite region contains a table surrounded by chairs. What it is for a table to be surrounded by chairs is vague (and it is a vague matter when molecules compose a chair or a table, for that matter). Nevertheless, when the demon is sorting through complex descriptions of regions of space to find the right one, her informant's advice is useful. Empty regions, for example, are out, as are uniformly packed regions, I suppose. If the Laplacian demon is lucky enough, she may even be able to pick out the only eligible candidate out of a pile of descriptions she is given – she may not be able to be sure her candidate counts (e.g. because that may require information about what went on outside the region, if e.g. origins or creators' intentions are essential to whether something counts as a table), but her informant's information is definitely some guide.

What the demon should *not* do, if she is given the task of finding a region in which a table is surrounded by chairs, is reason in the following way: "The distinction between regions where a table is surrounded by chairs and regions where there are no such things

is a vague distinction. However, the distinction I am called upon to locate is one that is a precise one – I have to pick one of a range of precise alternatives. My distinction, therefore, "unless it is vague, cannot fit the intuitive desideratum". So I had better ignore my informant's advice, and make my choice on other grounds". That would be a recipe for disaster.

Most of our metaphysical intuitions and metaphysically relevant deliverances of our theories of the world will be cast, in the first instance, in vague language. If Lewis's style of argument shows that we cannot, or should not, try to respect any of these constraints when coming up with a metaphysical theory, it proves too much. Is there something special about the case of composition, though? To explore this question, I think it would be useful to examine a pair of arguments offered in Sider 2001. The first is Sider's reconstruction of Lewis's argument (or a close relative) – it is presented more explicitly than Lewis's, so it will be easier to what, if anything, is supposed to be distinctive about composition here. The second of the pair is Sider's argument for temporal parts from vagueness. (An argument Sider classifies as "one of the most powerful" for temporal parts). I think Sider's argument from vagueness for temporal parts has the same problem that Lewis's argument has (though no doubt there are other ways to challenge it too, since it involves principles about time that are the subject of a great deal of controversy). Once it is clearer what the fault I have in mind is, one line of response for the endurantist to Sider's argument will become clear.

3. Sider's Argument for Unrestricted Composition

Sider offers us an argument that he describes as Lewis's, but developed in Sider's own way (p 120). Sider's main argument is valid (or can be made so without much difficulty), so if I am to say that the conclusion is not established I had better fault one of the premises. Sider's argument for unrestricted composition proceeds as follows (the following are taken from Sider 2001 p 123-125):

> P1: If not every class has a fusion, then there must be a pair of cases connected by a continuous series such that in one, composition occurs, but in the other, composition does not occur.

> P2: In no continuous series is there a sharp cut-off in whether composition occurs.
>
> P3: In any case of composition, either composition definitely occurs or composition definitely does not occur.

His conclusion is that every class of objects has a fusion.

By a "case of composition" (or "case" for short), Sider means nothing stronger than "a possible situation involving a class of objects having certain properties and standing in certain relations" (p 122). The argument as presented is not formally valid – to make it formally valid, one could supplement it with a principle about how continuous series can have ends which are different with respect to some on/off feature – e.g. either by having a sharp cut-off, or by having steps which are not determinately like the first step or not determinately like the last step, for example. (That is, they have sharp cut-offs or vague transitions from one state to the other). Perhaps we could hope to extract some such principle from the meaning of expressions in the premises – "sharp cut-off" for example.[4] Or perhaps we could add it as an independent premise. However we are to justify such a principle, let us take some such principle as read, and so allow that the modified argument is valid.

Sider should also restricts his discussion to composition of physical objects, or have some equivalent restriction. Someone who accepts restricted composition because they believe regions of space never form composites with physical objects, for example, or who thinks that propositions and puppies never form fusions, is not obviously going to be affected by an argument like Sider's, since they will have no obvious reason to accept the first premise (you can't get from propositions to puppies via a continuous series of plausible stages, or at least I do not see how it is to be done). Lewis suggests such a restriction for his argument (p 212 ftnt 9), and Sider would be well advised to take this restriction on board. I will assume on his behalf that he has.

[4] Much will also depend on what, exactly, a "sharp cut-off" is. Given the reference to determinacy in P3, I expect a sharp cut-off is intended to be a step in the series such that there is a transition from something being determinately one way to the next step being determinately the other.

So the first premise says that we can find a possible sequence of cases that stand at endpoints of a "continuous series", where each member of the series varies to an extremely small degree from the cases near it in the sequence.[5] We could think, for example, of a piece of granite in the first case, and the case at the other end being a case where the atoms that make up the granite are scattered across the galaxy – we can connect these quite disparate cases with a continuous series, for example by considering possible cases where at each step the atoms are slightly further apart. If P1 is correct, then if unrestricted composition is false, there is *some* such sequence of possibilities. (The exact case may vary depending on which principle of restricted composition is considered). I can imagine that P1 could be denied – for example, if someone thought that objects came in two radically different, precisely demarcated, kinds, and composition only occurred within a kind. Another basis for denying it is if one denied that objects *ever* were parts of other objects (see Rosen and Dorr 2002). Sider argues against this "mereological nihilism" on pp 176-180 of Sider 2001. But I am happy to grant P1 to Sider, at least for the sake of the argument, since I think most plausible restrictions on composition of the relevant sort will join him in accepting P1.

P3 is the premise that requires wheeling out the semantic conception of vagueness, and it is the premise to which Sider devotes the most attention. I expect it is the place where the most controversy will be focussed (and indeed, when Peter van Inwagen rejects Lewis's argument in van Inwagen 1990 chapter 18, he in effect does it by rejecting P3). Sider argues that if composition is vague, then there will be vagueness about how many non-abstract objects there are (expressed as a claim in the language of quantifiers and identity, with an entirely unrestricted quantifier, and a predicate for concrete objects that is not relevantly vague). An unrestricted quantifier, the identity symbol, and some sort of concreteness predicate have no room for multiple precisifications, Sider goes on to argue, so there is no way it can be a vague matter whether such a sentence is true. For what it is worth, I do not think Sider's defence here is compelling – one can agree with his claims

[5] Sider talks as if there will be "adjacent" members of a continuous series, but this may be a simplification on his part – sometimes continuous series may also be *dense*, so that between any two stages there is another (consider the rational numbers – between any two rational numbers you can find others).

about logical terms, unrestricted quantification, and his claims about vagueness in language, and still resist the ultimate conclusion that there is no vagueness in composition (see Nolan, in preparation. Andre Gallois has also pointed to trouble in Sider's argument for P3, in unpublished work). But for present purposes, I propose to grant P3 to Sider, if only to avoid the suspicion that our difference turns on a contention about the nature of vagueness.

Granting P1, P3, and the validity of the argument, at this point it is unsurprising that my problems will be with P2. Why would one accept P2? One might accept it if one accepted vague composition across the board, so that there was never a sharp cut-off. (This is less plausible than it might appear. There is no fusion where there are no objects, but we can get from some objects to no objects by a possible continuous process of subtraction of matter.) Or one might accept it if one antecendently accepted one of the extreme principles of composition that did not allow for cut-offs at *all* in composition, sharp or otherwise. One such principle is mereological nihilism, according to which there are never fusions (or other-than-atomic fusions), and another is unrestricted composition, according to which every case (excluding perhaps the empty case) has a fusion. But suppose you were inclined to think that sometimes there was a sharp, non-vague line in a continuous series of the sort covered by P1. Why would you accept P2? Here is what Sider has to say in defence of P2:

> To postulate such a sharp cut-off would be to admit that the realm of the macroscopic is in some sense 'autonomous' of the microscopic. By 'autonomous' I do not mean 'non-supervenient', since accepting a sharp cut-off in a continuous series of cases of composition does not threaten supervenience. Rather I mean that there would seem to be something 'metaphysically arbitrary' about a sharp cut-off in a continuous series of cases of composition. Why is the cut-off here, rather than there? Granted, everyone must admit *some* metaphysically 'brute' facts, and it is a hard question why one brute fact seems more or less plausible than another. Nevertheless, *this* brute fact seems particularly hard to stomach. (Sider 2001 p 124)

There's a fashion for throwing around the charge of "arbitrariness", but I am suspicious of it. There are several things a charge of arbitrariness could amount to, I take it, but

most are not applicable, and none are applicable uncontroversially. One thing it might be is a complaint about a lack of simplicity – sometimes a position can seem arbitrary or *ad hoc* if it is not the simplest or most straightforward alternative available. Someone who postulates some cut-off between cases, so that composition occurs in one but does not occur in the other, has a less simple theory than someone who needs no such division (e.g. a defender of Unrestricted Composition). I think the simplicity argument for unrestricted composition does indeed have something to be said for it – but it seems to me a quite different argument from one that might be extracted from considerations about vagueness. (It would also be better presented, I think, in its own right, not embedded in the middle of other principles that are superfluous to the simplicity argument).

The *postulation* of a particular cut-off may be arbitrary if there is no more reason for that cut-off rather than others relevantly like it. If I say that atoms have a fusion provided they are no more than 1.34 millimetres apart, you might legitimately reproach me for being arbitrary – why 1.34 mm, rather than 0.8 mm or 1.36 mm? Perhaps this is what Sider has in mind. This would, I think, misunderstand the position he is criticising – his opponent need only say that there is some cut-off, and not necessarily take themselves to tell exactly where the cut-off is. (Presumably they have rough-and-ready methods, but these do not need to be exhaustive, I presume). In any case, I do not think this is the kind of arbitrariness Sider is complaining about – his expression "metaphysically arbitrary", and his complaining about a metaphysically brute fact, suggests that he thinks there is something wrong with the view that there is some cut-off or other, and not merely with a view that stipulated that the cut-off was in a particular place without some evidence to locate it there rather than elsewhere.

We could interpret Sider as charging that principles about the extent of composition should not be primitive, but should be explainable in terms of something else more fundamental – but as Markosian 1998 points out and Sider seems to agree, mereological

principles do not usually receive a further analysis in terms of something else.[6] Markosian uses the language of there being "brute facts" about composition, and perhaps Sider is echoing this when he charges that "*this* brute fact seems particularly hard to stomach". Even if he is, though, it is hard to believe that it is the *bruteness* of a fundamental fact about composition that Sider is objecting to, since unrestricted composition (or any other answer to the composition question) would plausibly be just as brute.

Maybe Sider is asking for another sort of explanation of the cut-off (causal? teleological?), without which the cut-off would be "brute" but that's rather unlikely – it's not as if anyone in this debate thinks the facts about mereology require a story of their genesis, or that we need to "justify the ways of God to man" in explaining why He decreed that some things shall make up wholes and others not (or at least if *that* was the lack that someone was complaining about by calling something 'metaphysically arbitrary', we should not be concerned, if only because there is not a god of mereology who imposed those principles). So the supposed "arbitrariness" or "bruteness" is probably not meant to be something produced by the lack of a causal or teleological explanation. Unrestricted Composition seems to be as much a "brute fact" (unanalysed, not given a causal or teleological explanation) as the kind of fact Sider criticises here.

Sider uses the language of something's being "metaphysically arbitrary" in one other place that I could locate: on p 131 of Sider 2001. There, it is being used as an antonym of "metaphysically privileged" – a boundary is metaphysically arbitrary if there is nature does not "carve at the joints better than the rest" at that boundary, if there is no significant difference in the application of some eligible, natural metaphysical feature of the world. I suspect Sider is using the expression on p 131 differently from the way that he is using it in his defence of P2. After all, there's nothing here to suggest the defender of restricted composition thinks that composition is an unnatural gerrymander. The difference between a case where there is composition and a case where there is not is precisely a

[6] At the very least Sider 2001 does not offer any analysis of composition in terms of anything else. He does flirt with a theory that may provide some explanation, the "industrial strength composition as identity" view discussed on pp 159-161, though in the end he apparently rejects it.

difference where there is a significant fundamental metaphysical difference – facts about composition are natural, eligible metaphysical facts. At the very least, there is nothing to stop a defender of restricted composition from saying so. On the other hand, if we take the similarity of phrasing seriously and take Sider to be charging that a sharp cut-off would not correspond to a joint in nature or a natural metaphysical distinction, then Sider is mistaken about what the defender of restricted composition is committed to. The boundary can well be "metaphysically privileged" and not at all "metaphysically arbitrary" in the sense Sider 2001 p 131.

Insofar as my understanding of what a charge of "metaphysical arbitrariness" goes, either there is some sort of simplicity consideration being gestured at, or the position criticised is not guilty as charged. Perhaps there is something else that Sider intends, but if so I cannot see what it might be.

Sider's "argument from vagueness" fails to be compelling, even if we grant him all he wants in the theory of vagueness (especially P3 and the supporting arguments for it). While I have criticised Sider's argument for P2, I have not yet said very much about how one might motivate a denial of P2. I will return to this unfinished business in the context of outlining a positive proposal for the defenders of restricted composition in section 5. Sider's argument for temporal parts is similar to his argument for unrestricted composition – let me briefly discuss this argument.

4. Sider's Argument for Temporal Parts

Sider's argument for temporal parts is careful, and as a result somewhat technical. Permit me, then, a slightly unusual form of presentation. I will present a simplified version of Sider's argument, which will be enough for me to make it clear where the argument fails to convince. I will then present Sider's argument in full-dress form, but with little comment, in its own subsection. Those interested in the technicalities can then check for themselves whether what I say is fair to the structure of the argument, and those prepared to take my word for it will be spared.

Objects with parts come into existence and go out of existence. There is often vagueness about when they start or finish (at what microsecond does a statue begin to exist, or a car rust so badly it goes out of existence?) Let us suppose we have an object *O* with a precise time it comes into existence and a precise time that it goes out of existence, and let us suppose that it exists continuously (i.e. it has no gaps in its history). We may not have any very ordinary words for such an object, but if vagueness is only in thought and language, it is plausible that there will be some objects with precise, non-vague, lifetimes. Call the exact spatial and temporal interval that *O* exists in its Lifetime. Some intervals of time will be Lifetimes for objects, but Sider's opponent will think that sometimes we can define a period of time and pick out a succession of (one or more) objects, so that while a member of that succession exists at each instant of the time period, nevertheless the time period is not a Lifetime for any single object that overlaps one of the succession at each instant – there is no object that comes into existence at the start of the interval and goes out of existence at the end of the interval, while overlapping one of the succession of objects at each of the instants in the interval.

Let me illustrate with an example. Suppose we consider ten minutes of my life (say, noon to 12.10 pm on the first of January 2000). I exist at each of the times in that interval. A believer in temporal parts will think that there are instantaneous "stages" of me that exist at each instant in that interval, and will usually believe that there is some ten-minute stage of me that exactly fills that interval. That interval, together with me, according to a temporal parts theorist, specifies a Lifetime for an object (the ten-minute slice of me). Opponents of temporal parts will not necessarily agree that there is an object that overlaps with me for those ten minutes, but comes into existence at 12.00 pm and goes out of existence at 12.10. They may well deny that there is an object that overlaps with me for exactly that ten minutes, and that specifying me at each of the instants in that period picks out the location of a ten-minute long object.

Another example of a succession of longer-lived objects plus a period of time can illustrate another way there can be a failure to have a single object with a single Lifetime associated with the period plus objects. Take the ten minutes of my existence between

12.00 and 12.10 of 1/1/2000, plus all of Tony Blair's existence minus that ten minutes. (Assume for the sake of the example Blair have precise starting and finishing times – which may not be true, but it may be hard to specify objects that are both precise and familiar). Now, there is something that comes into existence at the start of the period specified, and goes out of existence at the end – Tony Blair. But still, that period of time does not specify a Lifetime relative to the objects specified – for there is no object that begins existence entirely overlapping Blair, hops into me for ten minutes in 2000, and goes back to overlapping Blair for the rest of its existence. Or so an opponent of temporal parts may well claim. (Indeed, some friends of temporal parts might not like this gerrymandered object – a defender of temporal parts does not need to endorse unrestricted composition across times).

Sider's argument, roughly, is that any specification of intervals and objects which exist at each of those intervals will succeed in defining a Lifetime. (Sider's actual argument is more general, and more careful, but this will do to display the general idea). Suppose not every specification of an interval of time plus objects at each of those times picks out a single object and its Lifetime. Then, there will be a continuous series of cases from a case where a Lifetime is specified, to a case where a Lifetime is not. For example: if we suppose that Tony Blair has precise temporal boundaries (which he might not, but maybe some precisification of "Tony Blair" does pick out something with precise boundaries), then we can get from that Lifetime to the specification of Tony Blair for ten minutes by a continuous series of very slightly shorter intervals.

Sider's second premise is that no such continuous series has a sharp cut-off between when a Lifetime is specified and when it is not – if the ten-minute interval does not pick out a Lifetime of a ten-minute slice of Tony, then there is not some sharp line where Lifetimes stop getting picked out between the full-sized Lifetime and the ten-minute interval.

Sider's third premise is that for any precise interval (and objects picked out precisely associated with that interval), either it definitely picks out a Lifetime for some object or it

definitely does not – there is no occupied space-time region that is only vaguely a Lifetime for an object. If there was this vagueness, Sider argues, then there would be vagueness about how many concrete objects there are at a given precise instant. But claims about how many (concrete) objects there are at an instant can be stated in the language of unrestricted quantifiers, identity, a non-vague specification of "concrete" and a precise time restriction (or some other non-vague piece of represntation of tense). So there cannot be vagueness about how many objects there are at a time.

The conclusion of these three premises is that *every* specification of an occupied spatio-temporal region specifies a Lifetime. This includes instantaneous specifications, so for any instant where I exist, there is an object that is a part of me that only exists for that instant – so there are temporal parts.

The argument is obviously analogous to Sider's previous argument. Again, I am happy to spot him the premise about a continuous series, for present purposes at least, and am prepared to allow him whatever principle he needs for the argument to be valid. Let me even allow him the principle about there being no vagueness in what exists at a time (though as well as any worries one might have about his justification or the theory of vagueness underpinning it, I think there are also disputable assumptions about temporal location). The problem this paper deals with is again to be located in the second premise. Why not have a precise cut-off between the occupied intervals that count as Lifetimes and the ones that do not? Sider allows that some versions of mereological essentialism may motivate this here (p 135), but offers arguments against those views later in his book (p 180-187). Apart from that, all he says is that sharp cut-offs would be "implausible" (p 135) – presumably because they are "metaphysically arbitrary". This part of the argument is independent of motivations from the "linguistic theory of vagueness", and as we saw it is difficult to see how it is supposed to have force for someone inclined to defend an alternative to Sider's temporal parts view.

The Full-Dress Argument

Sider's full-dress argument is more general, and it is more careful to ensure that the objects associated with the equivalent of "Lifetimes" are parts of the longer-lived objects used to specify them. The premises below are direct quotes of Sider 2001 p 134. The definitions of "assignment" and "D-fusion" (or diachronic fusion) are direct quotes from p 133. The definition of a minimal D-fusion is a very slight paraphrase of the one given on p 133, and the definition of fusion-at-t is a slightly reworked, though equivalent, version of the definition given by Sider on p 58.

Definitions:
fusion-at-t: *y* is a fusion-at-*t* of the members of a set *S* iff all the members of *S* are parts-at-*t* of *y*, and every part of *y* has a common part-at-*t* with some member of *S*.
(part-at-t is either a primitive or defined in terms of temporal parts, depending on the background theory).

assignment: any (possibly partial) function that takes one or more times as arguments and assigns non-empty classes of objects that exist at those times as arguments.

x is a *D-fusion* (diachronic fusion) of an assignment *f* iff for every *t* in *f*'s domain, *x* is a fusion-at-*t* of *f(t)*.

A *minimal* D-fusion is a D-fusion that exists only at the times in its assignment's domain.

Premises:
P1': If not every assignment has a minimal D-fusion, then there must be a pair of cases connected by a 'continuous series' such that in one, minimal D-fusion occurs, but in the other, minimal D-fusion does not occur.
P2': In no continuous series is there a sharp cut-off in whether minimal D-fusion occurs.
P3': In any case of minimal D-fusion, either minimal D-fusion definitely occurs, or minimal D-fusion definitely does not occur.

So the conclusion is that any objects at-a-time always form a minimal D-fusion at any time when they exist. A limit case of this means that at any instant an object exists, it entirely overlaps-at-that-instant an instantaneous object – so it has a temporal part at that time.

This statement of the final stages of the argument is my paraphrase, rather than a quotation from Sider, but I am prepared to accept for the sake of the argument here that his conclusion follows from his premises (though strictly we may wish to have a supplement analogous to the one suggested for the argument presented in section 3).

5. Return to Lewis – How to Link the Vague and the Non-Vague

Vague intuitions can constrain what we are to believe about non-vague matters, or so we believe. Section 2 suggested that if this were not so, we would be in serious trouble when trying to decide what to believe about fundamental metaphysics, since most of our intuitions and other evidence are represented vaguely. However, we can allow our views, even about mereological composition, to be guided by vague constraints, even if we adopt a view of vagueness according to which it is entirely a matter of "semantic indecision".

To see how this could work, let me return to Lewis's argument. As reconstructed, the central argument was

(1) The intuitive criteria for restricting composition are vague, but
(2) When composition occurs cannot be a vague matter, so
(3) When composition occurs cannot fit the intuitive criteria.

As I proposed, allow me to grant him both premises, for the sake of the argument. The problem is that Lewis's conclusion does not follow from (1) and (2). Even if composition is not vague, it can still fit vague criteria.

How can a phenomenon describable in non-vague terms satisfy a vague criterion? Let us first consider, by analogy, the case of a vague necessary condition. Suppose there is a custom in a small country that its dictator grants, once a year, tax-free status to one of the nation's citizens. Let us suppose that on the day of the announcement there is no relevant vagueness at the time about who counts as the citizens of the country – there are no people being born, or on the cusp of death on that day, there are no disputable cases of immigration or emigration, and so on. Let us suppose further that the dictator can choose a citizen in a way that will not have any indeterminacy due to vagueness – he will announce the lucky citizen's name clearly, and the citizens do not share names or have

names that sound too similar. Let us suppose, in fact, that in all relevant respects it will not be vague which person the dictator selects for this privilege.

Suppose now that it is also the custom that the dictator will select a bald citizen for this tax-free status. Can the dictator's non-vague choice satisfy this condition? Clearly it can. There will be borderline cases of baldness in the country, let us suppose – cases that on one resolution of the semantic indeterminacy in "bald" would turn out to be in the extension of "bald", and on some other resolution, not. There can also be citizens that would be bald no matter how one made semantic decisions about "bald" so far left unmade, and some hirsute citizens who are already determinately not-bald. If the dictator chose one of the borderline baldies, it would be vague whether the dictator had met the criterion. But if the dictator picked someone with a hairless head for the honour (as he might, given the rash of head-shavings that we can imagine in the lead-up to the announcement), he will have satisfied a vague necessary condition with a non-vague selection.

A similar point can be made with respect to sufficient conditions. Suppose the dictator also wants to ensure that a pension will be provided for every citizen who is "elderly" – it is vague at what age someone becomes elderly (and whether it just depends on age), and let us suppose the full spectrum of ages are represented in his state. Suppose his minister for social services is instructed to draw up a list of citizens to be given the first instalment of the pension at noon. Is there a way the minister can draw up a list that is not vague in any relevant way, but ensures that no elderly citizen is left out? Yes. A lazy minister could put every citizen on the roll – that will not leave out any elderly citizens! Or a slightly more responsible minister may decide to list everyone 15 years of age or over. This minister will leave some citizens off the roll, but will not be in any danger of it being vague whether or not she enrolled every elderly citizen. Non-vague divisions can satisfy vague criteria, both necessary criteria and sufficient criteria.

The possibility of composition satisfying a raft of vague criteria should now be obvious. Plausible criteria will probably involve a range of factors and often be domain specific,

so I will illustrate with just two examples. Suppose we say that a set of objects do not have a fusion if they are *all very far away from each other*. This is a vague criterion. A restricted (non-trivial) principle of composition could satisfy this, even if it could be stated in entirely non-vague terms. "all very far away from each other" is an expression with many precisifications[7], and if it is to be determinately true that every set of objects meeting this condition lacks a fusion, it had better turn out that all the sets that have fusions fall outside the extension of *any* of the precisifications of the expression "all very far away from each other". But this can happen – and if it does, we have a condition stated in vague language which is nevertheless a sufficient condition for there being no associated fusion.

If a necessary condition is met on all precisifications, it is determinately satisfied despite its vagueness. Likewise for a sufficient condition. Precise boundaries can meet all the precisifications of such conditions. If our intuitions come in the form of claims that in various circumstances (or in various paradigm cases) certain sets of objects do have fusions; and in certain other cases sets of objects lack fusions, then the fact that these intuitions are vague does not yet show that they cannot be simultaneously met by a non-vague principle of composition. Of course, it is possible to have a set of conditions, some necessary and some sufficient, that cannot be determinately satisfied together. The simplest case of this is when such conditions are mutually inconsistent, or have a partly inconsistent application – if I say a necessary condition to win a competition is to be over six feet tall, and a sufficient condition is that one is under five feet, then if a 4'6" person enters the competition, I cannot consistently pick a winner. Another way for the conditions to conflict is if their precisifications overlap in the right way – it may be that the precisifications of "tallish" and the precisifications of "shortish" overlap, even if nobody is determinately tallish and shortish: and if the dictator insists that all of his soldiers be tallish and some be shortish, then his poor minister of the army has been given an impossible job.

[7] Talk of "precisifications" suggests supervaluationism, of course, but the point here can be paraphrased into the vocabulary of other linguistic theories of vagueness.

Sometimes, interestingly, vague criteria will not be able to be satisfied by a non-vague boundary, even if they are satisfiable by a vague distinction. Suppose I say that every "small" heap of sand has a fusion, but no "non-small" heap of sand has a fusion, and suppose we have a continuous series of heaps, each differing by an individual grain of sand. There is no non-vague boundary I could draw that would determinately satisfy both conditions – if I exclude some determinately small heap, or include some determinately non-small heap, I have failed, and if I draw a line between two heaps that are small on some precisifications, and non-small on others, then I have not determinately satisfied either condition. I can, however, provide a vague description that will capture the small heaps and exclude the non-small heaps on every precisification - the most obvious one being "the small heaps are in, and the non-small heaps are out". If the intuitive constraints on composition are like these, then it will be impossible to satisfy them with a non-vague principle of composition.

I suspect people have been led astray here partly because of the form of intuitive answers to questions about composition that they have been looking for. For example, van Inwagen 1990, and following him Markosian 1998 think the defender of restricted composition has as an important focus the "special composition question". In Markosian this is framed as

> What necessary and jointly sufficient conditions must any *x*s satisfy in order for it to be the case that there is an object composed by those *x*s? (Markosian 1998 p 212)

Sider seems to attribute to Lewis the view that the way to provide a theory of restricted composition in accord with our vague intuitions would be to supply something vague in the schema "A class, S, has a fusion if and only if——" (Sider 2001 p 121-122), though I am more dubious that this assumption is in Lewis's argument.

If what intuition offered were statements of the form "A class S, has a fusion if and only if...", where the class S was specifiable in a non-vague manner, then Lewis and Sider would be right that given their conclusions about vagueness (Lewis's (2) and Sider's P3 and P3'), the answers offered by intuition would have to be rejected. Whether or not a class has a fusion is not a vague matter, according to Lewis and Sider – and so if it

obtained iff some vague condition obtained, there would be paradox looming. For in a case where it was indeterminate whether the condition obtained, it would obtain on some precisifications but not on others. But since the biconditional would be determinately true, it would obtain on all precisifications. A non-vague sentence, to the effect that a certain class had a fusion, would then be true on some precisifications but false on others. But a non-vague sentence has the same truth-value on all precisifications of vagueness. *Reductio*. I speculate that it is something like this reasoning that drives Lewis, Sider, Markosian and van Inwagen in the theoretical directions they take. If you think the linguistic theory of vagueness shows that a vague answer to the special composition question cannot be given, and that the only answers supported by intuition are vague, then you have several ways to go. You can conclude that we should reject the answers supported by intuition, as Lewis, Sider, and Markosian do (though Markosian 1998, unlike Lewis and Sider, rejects unrestricted composition, preferring a precise special composition answer that he thinks does not satisfy our intuitions). van Inwagen, on the other hand, accepts a vague answer to the special composition question, and so rejects the linguistic theory of vagueness (van Inwagen 1990 p 228).[8]

Those who think we must choose between vague intuitions about composition and the linguistic theory of vagueness are mistaken. Vague criteria drawn from intuition (or other sources) need not be in the form of jointly necessary and sufficient conditions. If our vague criteria are, instead, a list of some necessary and some sufficient conditions, a non-vague answer to the special composition question can in principle satisfy them all. This will be true, for example, if our intuitions are not completely opinionated. If intuition tells us that some clear cases are in, and some clear cases are out, it may be silent about a significant range in the middle. If they are silent enough, then the principles may be able to be all satisfied on every precisification of each of those principles.

[8] While vague intuitions can support the postulation certain non-vague demarcations between cases of composition and cases without composition, this of course does not prevent them from supporting postulations of vague demarcations as well. My guess is van Inwagen will prefer to keep his vague principle of composition, and continue to reject the linguistic theory of vagueness, even were he to acknowledge that vague intuitions could support non-vague boundaries as well.

It might even be that some of our intuitions might be more opinionated – some might even be in the form of "S has a fusion iff...". Even if this is so, constructing a theory of composition that did not respect these more opinionated intuitions need not be a matter of ignoring intuition altogether. Lewis has taught us that in analysis, sometimes near enough is good enough (see e.g. Lewis 1972 p 253). Even if some of our intuitions needed to be rejected, that does not mean that they are all rejected.

Suppose, as a first pass, we came up with an "iff" vague characterisation of restricted composition. Following Lewis's suggestions (Lewis 1986 p 211), it might be

(L) A class S has a fusion if, and only if, its members contrast with their surroundings more than they do with each other, are adjacent, stick together, and act jointly.

This cannot be right, if the linguistic theory of vagueness is correct. Some possible cases will satisfy some precisifications of the right hand side and not others. However, it is easy to see how we could keep some of the spirit of this characterisation of composition. For example, suppose we fell back on the following two claims:

(+) A class S has a fusion if its members contrast with their surroundings *much* more than they do to each other, they are touching, they are hard to separate, and act jointly.

(-) A class S *does not* have a fusion if its members contrast with each other no more than they do with their surroundings, they are very widely scattered, move almost entirely independently, and do not act jointly.

These two claims leave much undetermined about when objects have a fusion, and even together they are much less strong than (L). But someone who thinks that (L) is intuitive can hardly claim that (+) and (-) are not intuitive, since they are obvious and natural weakenings of (L). (They are "natural" weakenings of (L) as opposed to, say, (L) disjoined with some arbitrary sentence, which would also be a weakening of (L) but

might command little intuitive appeal in its own right). (+) and (-) seem to leave enough slack so that some non-vague principle of composition could entail them both.

It may be that in the end we are entitled to judge that intuitions like (+) and (-) are outweighed by systematic considerations or theoretical advantages, or even other intuitions. But this reckoning will have to be done on its own terms – commonsense intuitive principles like (+) and (-) cannot simply be dismissed on the grounds that they are vague, even if we endorse Lewis's story about vagueness. And just as well – for if Lewis could dismiss vague intuitions about composition out-of-hand, arguments could be mounted that we could dismiss any vague opinion when it came to what the world was really like, to the extent that "what it is really like" can be captured in non-vague terms (or more realistically, to the extent that aspects of vagueness can be avoided for some specific purpose). There is no shortcut to establishing unrestricted composition or temporal parts simply through considerations of vagueness.[9]

Daniel Nolan
Departments of Philosophy
University of St Andrews
St Andrews, Fife, Scotland KY16 9AL
UK

References

Braun, D. and Sider, T. unpublished. "Vague, so Untrue". http://fas-philosophy.rutgers.edu/~sider/papers/vagueness.pdf
Fine, K. 1975. "Vagueness, Truth and Logic". *Synthese* 19: 265-300
van Fraassen, B.C. 1966. "Singular Terms, Truth-Value Gaps, and Free Logic". *Journal of Philosophy* 63.17: 481-495
van Fraassen, B.C. 2002. *The Empirical Stance*. Yale University Press, New Haven.
Graff, D. 2000. "Shifting Sands: An Interest-Relative Theory of Vagueness". *Philosophical Topics* 28(1): 45–81.
Hyde, D. 1997. "From Heaps of Gaps to Heaps of Gluts". *Mind* 106,424: 641-659
van Inwagen, P. 1990. *Material Beings*. Cornell University Press, Ithaca NY.

[9] Thanks to Katherine Hawley and to audiences at the Australian National University, the University of Bristol, the University of Queensland and the University of St Andrews for discussion and suggestions.

Lewis, D. 1970. "General Semantics", reprinted in Lewis 1983, pp 189-229
Lewis, D. 1972. "Psychophysical and Theoretical Identifications", reprinted in Lewis, D. 1999. *Papers in Metaphysics and Epistemology*. Cambridge Uni Press, Cambridge. pp 248-261
Lewis, D. 1979. "Scorekeeping in a Language Game", reprinted in Lewis 1983, pp 233-249
Lewis, D. 1983. *Philosophical Papers, Volume 1*. Oxford University Press, Oxford.
Lewis, D. 1986. *On the Plurality of Worlds*. Blackwell, Oxford
Markosian, N. 1992. "Brutal Composition". *Philosophical Studies* 92: 211-249
Nolan, in preparation. "Balls and All"
Rosen, G. and Dorr, C. 2002. "Composition as a Fiction". In Gale, R.M. (ed). *Blackwell Guide to Metaphysics*. Blackwell, Oxford
Sider, T. 1997. "Four-Dimensionalism". *Philosophical Review* 106: 197-231
Sider, T. 2001. *Four Dimensionalism*. Oxford University Press, Oxford.
Soames, S. 1999. *Understanding Truth*. Oxford University Press, Oxford.
Williamson, T. 1994. *Vagueness*. Routledge, London.