On the Properties of Composite Objects

Michael J. Duncan

November, 2018

There are, broadly speaking, two views about composition and composite objects, though one is far more popular than the other. The popular view, which I shall call the *orthodox view*, is that composite objects are not identical to their proper parts,[1] individually or collectively. The other view, known as *composition as identity*,[2] is that composite objects are identical to their proper parts, taken together.[3] Our focus will be on just one version of composition as identity—strong composition as identity—which states that the identity in question is numerical identity. I will have nothing to say about other forms of identity or about non-standard views of numerical identity.

Figure 1 illustrates the difference between the orthodox view and strong composition as identity (hereafter just "composition as identity"). The figure shows on the left a rectangle (top) composed of two squares (bottom), according to the orthodox view of composition. Assuming the squares themselves are not composite objects, there is a total of three objects on this view: the two squares and the rectangle. The figure shows on the right the same rectangle (or squares) according to composition as identity. There is either a total of two

[1] The term "proper part" captures what is generally meant by "part" in ordinary English. For example, your hand is a proper part of your body, but not a proper part of itself.

[2] This way of dividing things up is somewhat non-standard in that it does not classify as a version of composition as identity the view that composition is merely *like* identity (so-called *weak composition as identity*—see Yi, 1999). This, however, is appropriate in the present context. Our concern is with views on which composition is an identity relation, and not with views on which it is merely like one.

[3] Together, as opposed to individually. The claim is not that composite objects are identical to each of their proper parts, but to all of them collectively (just as they are composed of all of their proper parts collectively).

objects (the squares) or a total of one object (the rectangle), depending on how we count them.[4]

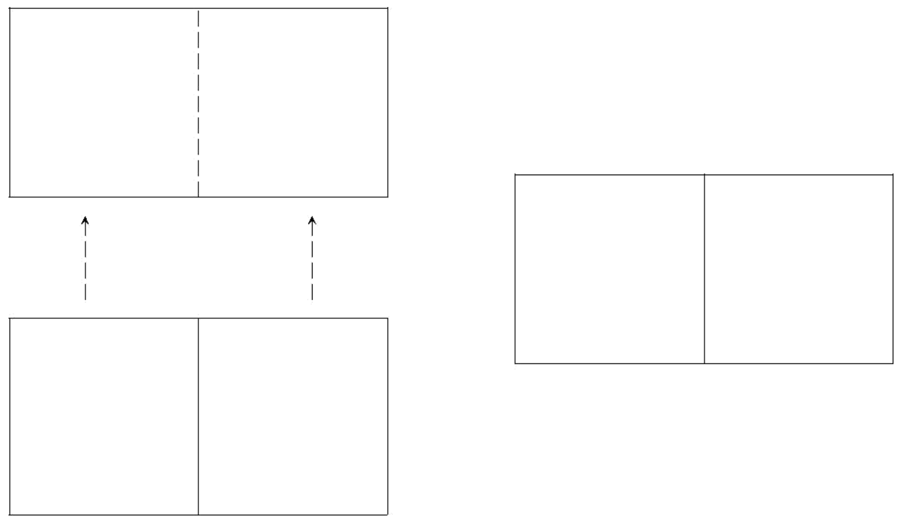

(a) The Orthodox View (b) Composition as Identity

Figure 1: A comparison of the orthodox view of composition and composition as identity. In subfigure (a) the rectangle at the top represents the composite object, the squares at the bottom its two halves. The dashed arrows indicate the relation of proper parthood. In subfigure (b) the rectangle and the two squares are identical.

Despite the popularity of the orthodox view, little has been said about the properties of composite objects and their proper parts if it is true.[5] This is a significant oversight. If the orthodox view is true, then we can ask separate questions about the properties of composite objects and the properties of their proper parts. Answering these questions is not as straightforward as it seems.

[4] It may well be consistent with composition as identity to say that only the count of one object is correct, and that each square is only half an object. What defenders of composition as identity cannot say is that there are three objects.

[5] See Van Inwagen (1990b, chapter 4), Parsons (2004), Sider (2007b), and Cameron (2014) for some exceptions.

I intend to defend three claims: (1) accepting the orthodox view requires a departure from the typical way in which we think about the properties of composite objects; (2) doing so leads to serious difficulties; (3) accepting composition as identity requires no such departure from the ordinary way in which we think about composite objects and their properties.

Whether these points should lead us to accept composition as identity is another matter. My own view is that they should; however, composition as identity faces problems of its own and I cannot hope to address those problems here.[6] As such, the main focus of the paper will be on the orthodox view, and on what we should say about the properties of composite objects if it is true.

I THE TROUBLE WITH THE ORTHODOX VIEW

Let us begin by thinking about the rectangle shown in Figure 1. On both views about composition the rectangle is a composite object, made up of two squares. What sort of properties does it have and how do these relate to the properties of the two squares? For instance, if the squares are grey what colour is the rectangle?

As Figure 2 illustrates, the answer is straightforward under composition as identity. On that view the question, "What colour is the rectangle?" has the same answer as the question, "What colour are the squares (collectively)?" Thus, if the squares are grey so too is the rectangle.

But what about under the orthodox view? The figure suggests that there is room for disagreement. Of course, the natural answer is that the rectangle must be grey also—but why couldn't the answer be "red"? The orthodox view seems to suggest that it should at least be

[6] For problems, see e.g., Yi (1999), Merricks (1999), Sider (2007b) and McDaniel (2008). For responses, see Wallace (2011a, 2011b), Bohn (2012), and Cotnoir (2013).

conceivable—if not possible—that the rectangle be red. But it is not. The proponent of the orthodox view has some explaining to do.[7]

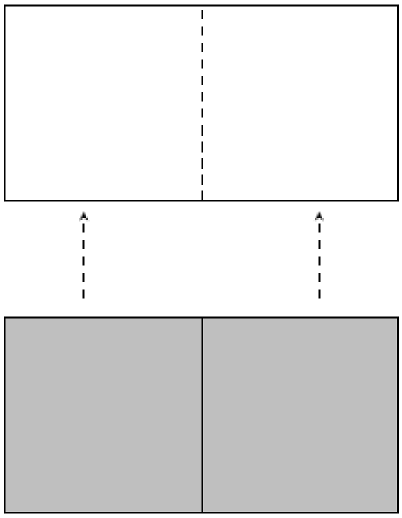

(a) The Orthodox View

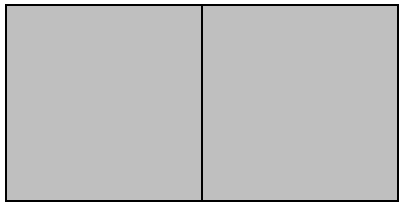

(b) Composition as Identity

Figure 2: A comparison of the orthodox view of composition and composition as identity. In subfigure (a) the rectangle at the top represents the composite object, the squares at the bottom its two halves. The dashed arrows indicate the relation of proper parthood. In subfigure (b) the rectangle and the two squares are identical.

We might also want to ask about the weights of composite objects. Take a chair, for example. Suppose that the proper parts of the chair (the legs, seat, back, etc.) collectively weigh x kilograms. How much does the chair weigh? The answer again, is obvious—x kilograms—but it is less obvious why that is the answer. And there is a further issue: if the chair weighs x kilograms, and the proper parts collectively weigh x kilograms, why don't these things, all together, weigh 2x kilograms? A typical answer, even amongst metaphysicians, is, "Because the legs, seat, back, and so on, are all *parts* of the chair." But

[7] See Cameron (2014) for an attempt to address some of these issues. In my view, his response is not successful, but I will not argue for that here.

this is really no answer at all. Why does the fact that these objects are parts matter? Without an answer to this second question, the original puzzle remains unsolved.

As before, composition as identity can solve it. Since the chair just is the legs, seat, back, etc., taken together, to sum together their weights would be to count the weight of the chair twice. I am not aware of any satisfactory response on behalf of the orthodoxy. Many these days would, I think, be inclined to say that the chair has its weight *in virtue of* the weights of its proper parts. Thus, the weight of the chair is not really anything over and above the collective weights of the chair-parts. But, as before, the question has not really been answered. Why does the fact that the chair has its weight in virtue of some fact about the chair-parts matter? If it has it, it has it, and it doesn't make a difference why.

We will return to this problem shortly; for now, let us go back to thinking about the rectangle. There is a deeper problem that needs to be addressed.

Suppose that the orthodox view is true and that one of the squares is black and the other white, as in Figure 3 below.

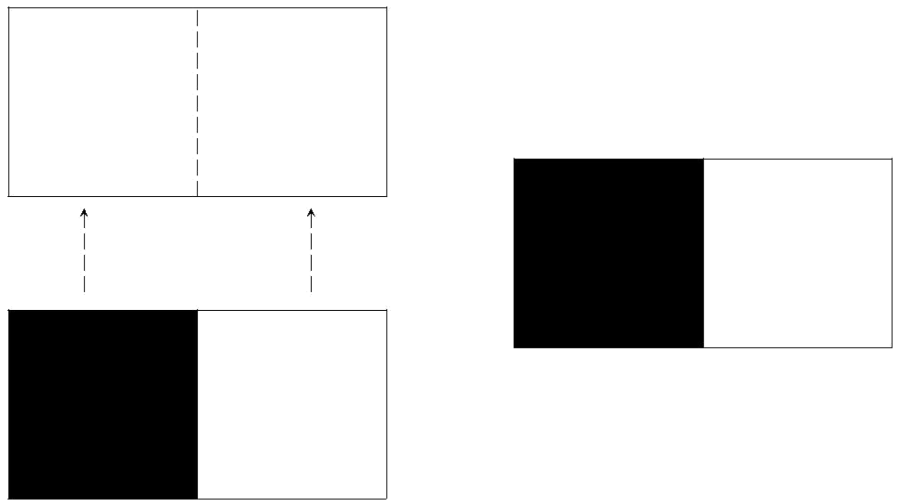

(a) The Orthodox View (b) Composition as Identity

Figure 3: A comparison of the orthodox view of composition and composition as identity. In subfigure (a) the rectangle at the top represents the composite object, the squares at the bottom its two halves. The dashed arrows indicate the relation of proper parthood. In subfigure (b) the rectangle and the two squares are identical.

What colour is the rectangle? The natural answer is “black and white”. But what exactly does that mean? Certainly, it does not mean the rectangle is black. That would entail that it is entirely black. Nor does it mean it is white, for the same reason. No. Intuitively, what it means is that the rectangle has a black proper part and a white proper part.[8] This is the answer that David Lewis (1986, p. 203ff) puts forward in his discussion of temporal parts, and most seem to accept it in that context. I suspect most will want to accept it here too.

Strikingly, the Lewisian answer is unworkable if the orthodox view is true. What we wanted to know was the colour of the rectangle; but what we were told was the colours of the squares. According to the orthodox view these are not the same. Thus, the Lewisian response completely fails to answer the question, “What colour is the rectangle?”

---

[8] More precisely, it seems to mean that the rectangle has one or more black proper parts, one or more white proper parts, and no proper parts of other colours. For ease of exposition I will continue to write, “a black proper part and a white proper part”.

There is a way that one could accept both the orthodox view and the Lewisian response without this consequence. For instance, one could say that the sentence, "The rectangle is black and white," does not in fact attribute any colour property to the rectangle at all. Instead, one could say that it is a claim about the colour properties of the squares. On such a view one could still say that the rectangle has a certain property—*being black and white*—but one could not say that this property is a colour property. For it is not. It is the property of having proper parts with certain colour properties. That is not the same thing.

This response also opens up a solution to the problem discussed earlier about the weights of composite objects. If, "The chair weighs x kilograms," actually attributes a weight to the chair's proper parts (collectively) and not the chair as it seems to, then the view says that the chair and the chair-parts together weigh x kilograms, as it should. After all, on this view the chair has no weight, at least strictly speaking.

The view, however, is very problematic for several reasons. First, it wrongly implies that the property *black and white* is not a colour property. Second, it wrongly implies that *black and white* is a relational property, and that a sentence attributing *black and white* to an object is not entirely about that object.[9] Third, it seems to imply either that (i) no composite object has colour properties, or that (ii) no composite object has what I will call *heterogeneous* colour properties—colour properties that vary across space (like *being black*

[9] Some might conclude from this that the view wrongly implies that *black and white* is not an intrinsic property either (e.g., Botterell, 2004). (Botterell makes his point in relation to temporal parts, but the point generalises to spatial parts as well.) Lewis (1983, p. 197), for instance, characterises an intrinsic property in the following way. "A sentence or statement or proposition that ascribes intrinsic properties to something is entirely about that thing." I am sympathetic to this characterisation of intrinsic properties, but others may not be since it implies that properties like *having a proper part which is black* are not intrinsic if the orthodox view is true. Note, however, that if composition as identity is true, the problem does not arise.

and white) —but that there are composite objects with *homogeneous* colour properties (like *being red*).[10]

Each of these possibilities is very strange. What would a composite object without colour properties even be like? The view seems to imply that we do not—and cannot—see many objects that we seem to see. My chair (or what I think is my chair) does not have a perfectly uniform colour. Thus, on either view, it has no colour properties. And therefore—presumably—I cannot see it.

And what if there *are* no non-composite objects at all—what if the world is "gunky"? What if there are no (visible) composite objects with genuinely homogeneous colour properties? This second possibility is not at all far-fetched. Closer inspection of pretty much any composite object that appears to have a uniform colour will reveal that it really does not. Something that appears to be perfectly white is unlikely to be so. According to the view we are considering, then, sentences like, "That thing is white," are almost always false, or are true but do not really attribute a colour property to the thing in question.

All this seems absurd. Proponents of the orthodox view should reject the Lewisian response. Clearly, then, some other account of the meaning of "black and white" is needed—one which does not make reference to the proper parts of the relevant object. Such an account is not as attainable as one might think. Our rectangle seems to instantiate both *black* and *white*, and yet these are incompatible properties. (Part of what it is to instantiate *black* is to not instantiate *white*.) Thus, the rectangle cannot be either black (simpliciter) or white (simpliciter)—for that

---

[10] Properties which vary across time are also heterogeneous, however since our focus is on the spatial variant, when I say "heterogeneous" I mean *spatially* heterogeneous.

Heterogeneous properties are what Parsons (2004) calls *non-uniform distributional properties* (or close to it). I avoid Parsons' terminology here because it is too fine-grained for our purposes. The distinction between heterogeneous properties and homogeneous properties which we will make use of is (at least roughly) his distinction between non-uniform distributional properties and the conjunction of uniform distributional properties and non-distributional properties.

seems contradictory—yet, somehow, it must instantiate both. How, if at all, can we account for all of this under the orthodox view?

II Alternative Accounts of Heterogeneous Properties

Luckily we have some resources to draw upon. If we are not allowed to appeal to the properties of a composite object's proper parts then we have the same problem as those trying to account for similar properties in objects which have no proper parts. Objects which have no proper parts despite being extended in space are known as *extended simples* and have been subject to investigation by philosophers (see e.g., Markosian, 1998; McDaniel, 2003; Simons, 2004; Markosian, 2004; Braddon-Mitchell & Miller, 2006; Tognazzini, 2006; McDaniel, 2007a, 2007b, 2009; Spencer, 2010; Jaeger, 2014). The challenge for those who think extended simples are possible is to explain how they could have properties that vary across space, like *black and white*, if they have no proper parts.[11] To say that an extended simple is "black and white" cannot be to say that it has a black proper part and a white proper part since, by definition, extended simples do not have proper parts.

We have seen that having proper parts is no help either, so we will need to avail ourselves of exactly the same resources to explain how composite objects can have such properties. This perhaps shouldn't come as a surprise. After all, a black and white extended simple seems to differ from a black and white extended composite only in terms of its relational properties. It seems that the two kinds of objects are black and white in exactly the same way.

Discussion of extended simples has yielded a range of strategies for accommodating qualitative variation across space. Let us discuss them one at a time.

[11] Assuming that they could have such properties. It is open to the defender of extended simples to deny that they can have heterogeneous properties. That, however, is a bold move.

Stuff

One strategy is to appeal to matter or "stuff". Ned Markosian (1998b, 2004) argues that when we say of an extended simple that it is "black and white" (for example), we are actually talking about properties of portions of matter or about portions of stuff.[12] While a black and white extended simple does not have a proper part which is black or a proper part which is white, according to Markosian it is constituted by some matter which does have a black proper part and a white proper part. Thus, according to Markosian, when we say that an extended simple is "black and white" we mean that it is constituted by some stuff which has a black proper part and a white proper part.

Markosian's view seems unappealing for various reasons. But, even setting those aside, it is no help to us here—it has the same deficiency as the Lewisian strategy. If we want to say that an extended simple *itself* has colour properties then we need to account for *its* properties. Appealing to facts about the colour properties of lumps of matter or stuff will not help us do this. Of course, as with the Lewisian response, one can always bite the bullet and accept that extended objects, composite or simple, do not have colour properties (or at least heterogeneous colour properties); but this seems a high price to pay.

As before, there are two such views. On one we say that composite objects have homogeneous colour properties but do not really have heterogeneous colour properties. On the other, we say that composites themselves lack both kinds of property. Neither option is very appealing. It is quite strange to think that when we say that an object is black we are talking about the object itself, but when we say it is black and white we are talking about a

[12] Stuff is the more general notion. Stuff need not be physical, and need not have mass whereas matter (at least arguably) does. In other words, matter is a kind of stuff.

portion of matter and not the object. And it is perhaps even stranger to think that objects (simple or composite) do not really have colour properties like *black and white* at all.

Relativised Properties

A more effective strategy is to appeal to spatially relativised properties. An oft-discussed approach to dealing with problems relating to persistence and change is to suggest that objects which change over time do so by having their properties relative to times. There are at least two variants of the view (see e.g., Lombard, 2003). First, we might say that the property *black* is really a relation, i.e., *black-at*, which holds between objects and times. Thus, an object might change properties over time by standing in the *black-at* relation to one time and the *white-at* relation to another. Second, we might say that the property *black* is really a time-indexed property: i.e., *black-at-*t_1. On this view a changing object simply has different time-indexed properties: e.g., *black-at-*t_1 and *white-at-*t_2. Both methods explain how something can be both black and white without contradiction.

The same two strategies can be adopted with regard to spatial rather than temporal variation. That is, we can say that a black and white object such as our rectangle stands in the *black-at* relation to one region of space, s_1, and the *white-at* relation to another region of space, s_2; or that it has the property *black-at-*s_1 and the property *black-at-*s_2.

This approach has at least two benefits. First, it can accommodate the intuition that the rectangle instantiates the properties *black* and *white* without being either black simpliciter or white simpliciter, since we can say that to instantiate the property *black* is to stand in the relation *black-at* or to have a space-indexed property of a certain kind. Second, it seems to fit quite well with how people talk about objects. For instance, it seems quite natural to say that the rectangle is “black *there* and white *here*”.

The view does, however, suffer several drawbacks. To begin with, properties like *black* and *white* simply don't seem to be relations, let alone relations to spatial regions . While it may not seem unintuitive that an object be black relative to one *time* and white relative to another it seems less intuitive that it be black relative to one *place* and white relative to another. Nor does a black object seem to instantiate a different property at different locations, as is entailed by the second variant of the view.

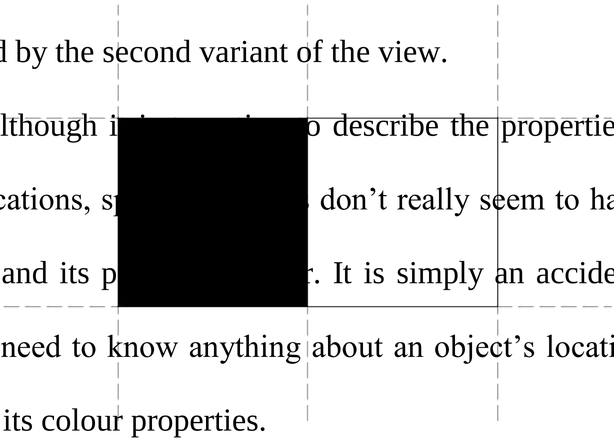

Furthermore, although i [illegible] o describe the properties of the rectangle with reference to spatial locations, s [illegible] don't really seem to have anything directly to do with the rectangle and its p [illegible] r. It is simply an accident that the rectangle is where it is. We don't need to know anything about an object's location to know everything there is to know about its colour properties.

Worse still, the relativisation approach implies that moving an object around in space changes its colour properties. Moving the rectangle from one location to another changes its colour properties from *black-at-*s_1 to *black-at-*s_2 (say), or changes which locations the object stands in the *black-at* relation to. An advocate of this approach can of course reply that the same property and object is involved throughout; but it certainly seems that *no* aspect of an object's colour properties need change when it is moved.

In addition to these problems the relativised property strategy faces what seems to me to be a devastating objection. Suppose our rectangle occupies a spatial region such that it is *black-at-*s_1 and *white-at-*s_2, as in Figure 4 below. If we flip the rectangle 180 degrees it will end up with the properties *white-at-*s_1 and *black-at-*s_2 (see Figure 5).

S_2

Figure 4: A rectangle located at the sum of the regions s_1 and s_2.

Or, on the other variant: the rectangle will go from standing in the *black-at* relation to s_1 to standing in the *black-at* relation to s_2. That is, its colour properties will change.

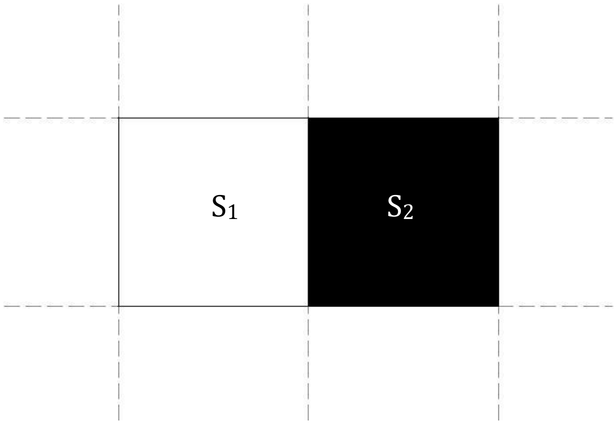

Figure 5: A rectangle located at the sum of the regions s_1 and s_2.

That the rotation results in a change in colour properties is not the full extent of the problem. Suppose that we had inverted the colour properties of the rectangle instead, by painting the black part white and the white part black, for instance. What then? We would have had the same result. The rectangle would have been *white-at-*s_1 and *black-at-*s_2. The relativised property strategy cannot distinguish between the two cases in the way we want. It can, of course, distinguish them in the terms we have laid out—in one case the rectangle is rotated; in the other the colours of its proper parts are changed. But it cannot distinguish them in terms of the colour properties of the rectangle itself. The rectangle ends up with the same colour properties either way. And that is the wrong result. Rotation and colour inversion should have opposite effects on the colour properties of a rectangle. The former should not change them whereas the latter should.

A similar problem arises if we imagine both rotating the rectangle and inverting its colour properties. On the view we are considering these actions result in no overall change in the rectangle's colour properties, despite the fact that inverting the colours should result in just such a change. Again, this is the wrong result.

The only escape appears to be to argue that we cannot make sense of inverting the colour pattern of the rectangle in the first place. That is a desperate move indeed. The relativised property view, then, seems to drive too great a wedge between an object and its properties. When we rotate the rectangle without changing its colour properties, what we want to say is the half of the rectangle which is black stays the same. However, this approach does not allow us to say anything of the sort. After all, a half of an object is a proper part of it. What we need to be able to say is that it is the rectangle itself, and not some proper part of it, that instantiates the property *black* (relative to a spatial location). Doing so, however, leaves us powerless to account for the difference between rotating the rectangle and inverting its colour properties.

Adverbialism

A close cousin of the relativised properties strategy is *adverbialism*.[13] Instead of relativising properties to spatial locations the approach is to relativise the instantiation relation itself. According to the second variant of the relativised properties strategy, an object instantiates the time-indexed property *black-at-*t_1; according to adverbialism the object bears the time-indexed relation *having-at-*t_1 to the property *black* (Van Inwagen, 1990a). In short, the idea is

[13] Advocates of the temporal variant of adverbialism include Johnston (1987), Lowe (1987), Haslanger (1989), and Van Inwagen (1990a). Van Inwagen (1990a) seems undecided between the relativised properties view and adverbialism (see pp. 247–248), so he can also be seen as a defender of the former.

that a changing object may be both black and white by instantiating *black* one way (“t_1-*ly*”, for example) and *white* another (“t_2-*ly*”, say). The properties *black* and *white* are incompatible, but only if they are instantiated in the same way.

On the spatial variant of adverbialism the strategy is to say that our black and white rectangle instantiates black s-ly and white s_2-ly (for spatial regions, s_1 and s_2).

Adverbialism shares most of its strengths and weaknesses with the relativised properties approach, though it does have some advantages. Foremost among these is that colour properties come out as intrinsic and non-relational. For instance, the view allows that our rectangle has blackness as an intrinsic property which is not relative to spatial location. It also allows us to say that an object instantiates *black* and also *white* without being either black or white simpliciter.

Again, however, this view seems more plausible in the temporal case. It seems less objectionable that an object have properties in different ways at different times than at different spatial locations. Perhaps, as an object moves through time it instantiates its colour properties differently. However, intuitively, an object’s location in space makes no difference to the way it instantiates its colour properties. The spatial variant of adverbialism implies that no two objects which occupy different spatial locations instantiate the same property in the same way. It allows that two people may have on the same colour shirt, but it does not allow that two people have on shirts which have the same colour *in the same way*. Yet surely my white shirt is white in just the same way as anyone else’s is.

And spatial adverbialism also faces the rotation problem. If we rotate our rectangle 180 degrees and invert its colour properties then, according to adverbialism, the rectangle has the same properties as it started with, and in the same way. Yet surely the colour properties of the rectangle have changed—in fact, we stipulated that they did.

Tropes

Given the failure of the previous strategies it looks like a different kind of approach might be in order. One such approach is to appeal to short-lived tropes (e.g., Ehring, 1997; McDaniel, 2009). Tropes are properties which are abstract particulars. In short, no two tropes are the same. Tropes may, however, resemble one another. The rough idea is this: an object is black if and only if it instantiates a blackness trope, where a blackness trope is one which falls into a class of (different) tropes which resemble each other in a certain way (McDaniel, 2009). Ehring (1997) proposes short-lived tropes as a solution to the problem of temporary intrinsics; McDaniel (2009) proposes the theory as a solution to the problem of qualitative heterogeneity in extended simples.

The solution works as follows. An object is black and white (across time or space) just if it instantiates a blackness trope at one location (in space or time) and a whiteness trope at another. In the case of the rectangle in Figure 4 , the trope view says that the rectangle instantiates a blackness trope which is located at s_1 and a whiteness trope which is instantiated at s_2.

The key feature of the view is that, on it, objects do not instantiate properties relative to different locations. Instead, objects exemplify tropes which are located at different locations. Both the instantiation relation and the property remain non-relativised—it is the locations of the tropes that is relative. This helps to avoid the problems associated with the previous two views. Properties are instantiated in the same way at different locations, they may be had intrinsically, and they need not be relational. Furthermore, we are able to say that the same object can instantiate both *black* and *white* without being either black or white simpliciter.

As with the previous two views, however, it falls foul of the rotation problem. If we rotate our rectangle 180 degrees then we change which kinds of trope are exemplified at which locations. Inverting the colour of the rectangle has exactly the same effect. There is nothing more to; and so we cannot distinguish the two states, even on the trope view.

The trope view also entails that no object exemplifies the same trope at two locations. Thus, strictly speaking, no two objects located at different regions of space have the same colour property. This is the same objection that the relativised property strategy faced. As with that view, the trope theorist can reply that having the same property merely requires exemplifying a trope of the same kind. In this way, she may perhaps claim to avoid the objection, though a view in which the very same property is instantiated will no doubt seem preferable to some.

Anti-Reductionism

The views we have discussed all have something in common: they all aim to reduce heterogeneous properties to homogeneous properties in some way. More specifically, they all offer an analysis of the concept BLACK AND WHITE in terms of the concepts BLACK and WHITE. As we have seen, however, doing this leads to serious difficulties (at least given the orthodox view of composition). It is worth considering something different.

We noted earlier that the rectangle seems to be neither black (simpliciter) nor white (simpliciter). Perhaps then we should say, following Parsons (2004) that the rectangle has some *other* property—namely, *black-and-white*—which is not reducible to the properties *black* and *white*.[14] If we do so we can avoid the contradiction that we started with. On this view no object is both black and white. To think this is to confuse an irreducibly heterogeneous property (*black-and-white*) for the conjunction of two homogeneous properties (*black* and *white*).

[14] I use "*black-and-white*" (rather than "*black and white*") where appropriate to indicate that I am talking about an irreducible property.

Appealing to irreducible heterogeneity is by far the best option among those we have discussed. It avoids the rotation problem. Consider Figure 5 again. On this view the rectangle in the figure instantiates the irreducible property *black-and-white*. What happens if we rotate it? It is tempting to say that the rectangle remains black-and-white, as we would expect (*cf.* McDaniel, 2009, p. 330). This is a step in the right direction, but, as with the other views, we get the same result if we invert the colour of the rectangle. That, of course, is not what we want.

Here, however, we can appeal to what we might call "directional" or "orientation" properties (*cf.* McDaniel, 2009, p. 330, fn. 13).[15] We can say that what changes when we rotate a black-and-white rectangle is its orientation, and not its colour property. On the other hand, when we invert the colour property of the rectangle, what we change is its colour property (from black-and-white to white-and-black), and not its orientation. More specifically, we can say that it goes from being black-and-white and being oriented "up" to being white-and-black and being oriented "up", whereas the rotated rectangle goes from being black-and-white and being rotated "up" to being black and white and being oriented "down".

The solution does not come without a price. Since we will want to say that it is the rectangle itself that rotates, and not just its proper parts, the solution seems to commit us to brute orientation properties. It is important to note, however, that this is a requirement of any theory which presupposes the orthodox view of composition. After all, there are two ways for a rectangle to occupy a single (rectangular) region of space; four ways for a square to occupy a single (square) region of space; and infinitely many ways for a disk to occupy a single

[15] McDaniel attributes the idea to David Lu. (Note that McDaniel discusses this point in relation to a somewhat different problem. The problem he discusses, and the rotation problem discussed here, are, nevertheless, closely connected. The discussion in this section is partly inspired by his discussion on page 330.)

(circular) region of space (assuming that space is continuous).[16] Thus, even proponents of the other accounts of heterogeneous properties discussed earlier must appeal to orientation properties. (Doing so won't help them with the rotation problem, however.)

The resulting view is a strange one. To start with, it is unclear whether we should attribute orientation properties to asymmetrical objects. And if we don't, why not? We might also wonder how to apply the idea to something like a rotating disk. It is natural to describe the rotation of a disk in terms of the relative (angular) velocities of its proper parts. When a disk rotates the proper parts towards the edge travel at a higher velocity that the proper parts towards the centre, for example. The rate of rotation of a disk can be characterised in these terms too. Under the distributional property account, however, we will need to say that the disk changes its orientation at a certain rate. Very well; but what if the disk is not perfectly rigid? Think of the inner half of the disk and the outer half. What if these change orientation at different rates? What then of the orientation of the whole? It looks as though we shall have to say that orientation properties can be heterogeneous too. The result: even more unwanted properties.

Finally, consider how orientation properties come to be instantiated. Suppose we construct a rectangle at region *R* by attaching together two squares. There are two ways the rectangle can be located at *R*: it can be located in the "up" direction or the "down" direction. But which is it? I can think of no reason to answer either way; however, it seems that it must be one or the other on this account. And what if we replace the rectangle at *R* with a different one? Is the new rectangle oriented the same way or some other way? Appealing to orientation properties raises more questions than it answers.

---

[16] I focus on two-dimensional objects for the sake of simplicity. Naturally, the same points apply for three-dimensional objects.

These are general problems which apply to the other views we have discussed. Anti-reductionism also faces problems of its own. The key feature of anti-reductionism is that heterogeneous properties like *black-and-white* are irreducible in the sense that the concept BLACK AND WHITE cannot be analysed in terms of the concepts BLACK and WHITE.[17] If they were reducible then we would be back where we started: we would have to say how an object can instantiate both *black* and *white* if these are incompatible properties.

This feature is also the source of numerous difficulties, all closely related. I will focus on just a few.[18] First, it simply doesn't seem true that the property *black and white* is irreducible. On the contrary, it seems closely related to the properties *black* and *white*. (It is no coincidence that the words "black" and "white" appear in the expression "black and white".) This explains why a black and white object seems to instantiate both *black* and *white*.

Second, consider various different objects made up of black and white proper parts in different arrangements. Each of these objects has a colour property which bears certain similarities to the others.[19] In addition to this, some of the resulting colour properties are more similar to each other than others. (Some have more black than white, for instance.) If each of these colour properties is irreducible, however, we cannot explain why some are more similar than others (Sider, 2007a; McDaniel, 2009), and why they are more similar to each other than they are to other heterogeneous colour properties such as *red and blue*. (Intuitively,

---

[17] Some characterise anti-reductionism as the view that heterogeneous properties are primitive, in the sense of being ungrounded or ontologically independent (e.g., McDaniel, 2009; Gilmore, 2014). This, I think, is a mistake since there is no reason why a composite object cannot be *black-and-white* and have this property in virtue of the fact that its proper parts are black and white. Such a property would be irreducible but not ungrounded.

[18] See Sider (2007a) and McDaniel (2009) for others. I do not discuss those difficulties for the sake of brevity, and also because I do not find all of them convincing.

[19] See McDaniel (2009 p. 329), Figures 3 and 4, for examples of what I have in mind.

the reason is that objects which are black and white have proper parts which instantiate *black* and *white* whereas objects which are red and blue have proper parts which instantiate *red* and *blue*.)

Finally, positing so many irreducible properties is far from parsimonious. On this view there is an irreducible colour property for virtually every imaginable distribution of colours in space. Furthermore, there are irreducible orientation properties, and irreducible similarity relations between the irreducible colour properties.

These are, I think, serious problems. They are not fatal, but they are troubling enough that I think we have reason to take composition as identity very seriously, despite its problems. Let us see how it fares in comparison to the orthodox view.

III Composition as Identity

If composition as identity is true the question, "What colour is the rectangle?" is essentially the same as the question, "What colour are the squares?" But we need to be careful when interpreting the second of the questions, for it is ambiguous. The question is not, "What colour is *each* of the squares?" but rather, "What colour are the squares *together*?" After all, the claim is not that the rectangle is identical to each of the squares—the claim is that it is identical to the squares *together*, just as it is composed of the squares together.

So what colour are the squares together? Certainly, they are not black. Only one of them is. Nor are they white. Again, only one of them is. Hence, it seems correct to say that they have some other property: *being black and white* or (even better) *being half black and half white*. So far this is in agreement with the anti-reductionist view and appears to be quite different to the Lewisian response. The difference between the current view and the anti-reductionist approach, however, is that on the current view the property *black and white* is reducible to the properties *black* and *white*. Saying that the squares are black and white is equivalent to saying that one of them is black and one of them is white.

Even opponents of composition as identity should accept this. By accepting composition as identity, however, we get the added benefit of being able to conclude the following: to say that the rectangle is black and white is (more or less) to say that one of the squares is black and the other white. And because the rectangle just is the squares, we can say this without having to say that the rectangle itself has no colour properties. The Lewisian response is therefore much more plausible if composition as identity is true.

Nevertheless, we still need to ask if the resulting view can meet the desiderata we came up with earlier. Can it accommodate the intuition that the rectangle instantiates *black* and *white* without being black or white simpliciter, for instance? It can. To see this we should return to thinking about the squares. One of the squares is black and the other white. Picture that in your mind. It should seem to you that both of the properties (*black* and *white*) are instantiated whenever the squares exist. Furthermore, this seems true in just the same way that both *black* and *white* are instantiated by the rectangle.

I will admit that it seems wrong to say that the squares *collectively* instantiate both *black* and *white*. But it also seems to me that the relevant intuition is only that *black* and *white* are instantiated *in some way* by the squares. Intuitively, *part* of the rectangle is black and *part* of the rectangle is white. Similarly, *one* of the squares is black and *one* of the squares is white. Facts about the squares guarantee that *black* and *white* are instantiated without it being the case that the squares are both black (simpliciter) and white (simpliciter). (The squares, collectively, are such that one of them is black and one white, just as the rectangle is such that part of it is black and part white.)

This solution seems to me to fit the facts remarkably well. The problem was to say how *black* and *white* could both be instantiated without contradiction, while also allowing that it is the composite object itself that instantiates them. If these colour properties are instantiated by the proper parts of the composite, and those proper parts are numerically distinct from the composite, then the fact that *black* and *white* are instantiated says nothing about the colour properties of the composite. If, on the other hand, the proper parts are identical to the composite, the fact that one is black and the other white *does* say something about the colour properties of the composite—it says that the composite is black and white.

It fits the facts about orientation remarkably well too. Take two squares, *a* and *b*, and put them together to form a rectangle at location *R*. If *a* is located at the left half of *R* (relative to some point), then the rectangle has one orientation; if *a* is located at the right half, then the rectangle has another orientation. However, by accepting composition as identity we open up the possibility of explaining this talk of orientation in terms of the relative locations of the squares. No irreducible directionality properties are needed. And—more importantly—we don't need to commit to any metaphysical distinction between these two orientations. Is the rectangle oriented "up" or "down" when *a* is on the left? We can say what we like. There is no further fact hidden away from view.

Not so under the orthodox view. Under the orthodox view it seems there must be an undiscovered (and likely undiscoverable) fact about whether the rectangle is oriented "up" or "down".

The problem facing all of the alternatives we have discussed was essentially that they could not distinguish one side of the rectangle from the other. Thus, the problem with the orthodox view seems to be that it makes parthood *external* to an object. What are called "proper parts" on that view are too far separated from the composite object itself. Composition as identity can do better; it imposes no such separation.

IV CONTINUOUS VARIATION

I want to conclude by addressing an objection. Parsons (2004) offers two scenarios designed to show that heterogeneous properties are not reducible to homogenous properties. Parsons' argument is intended as an attack on someone who endorses both the orthodox view of composition and the Lewisian view of heterogeneous properties. As such, it is not exactly targeted at the view I am endorsing. But one might still think it has some force against my position. It needs to be addressed.

Consider first an extended simple which is black and white. On our view, "x is black and white," entails that x has a black proper part and a white proper part. Thus, on our view such an extended simple is impossible. But, Parsons claims, extended simples are possible, as are extended simples with heterogeneous properties. So our view is false.

Second, consider an extended object which exhibits continuous qualitative variation across space. Parsons offers the example of a cloth with the full spectrum of colours ranging continuously from red (on the left) to blue (on the right). If we imagine a line running the full length of the cloth from left to right, it is as if a different (homogeneous) colour is instantiated at every point on the line.[20] Now, on our view it is as if a different colour is instantiated at every point on the line because there is a proper part of the cloth at every point which instantiates a different colour property. But, Parsons argues, it is possible for such a cloth to fail to have the necessary point-sized proper parts. Thus, the colour property had by the cloth cannot be reducible to homogeneous colour properties of its proper parts.[21]

---

[20] I say "as if" since Parsons denies that a different colour is instantiated at every point on the line.

[21] The argument can be made using a much simpler case. Consider an object which is as if it is white all over except at a single point at which it is black. The possibility of such an object, lacking point-sized proper parts, is sufficient for Parsons' argument to go through.

As with many arguments which rely on claims about metaphysical possibility, these are difficult to assess. Is an object exhibiting continuous variation in properties across space really possible? I am not sure.[22] But suppose that it is possible; I think the arguments are nevertheless not effective against the view I am defending.

In both instances, Parsons bases his claims about possibility on the fact that the objects and properties in question *seem* possible. He suggests that, absent evidence to the contrary, we should take things at face value. Rather than question this, I want to question whether extended black and white simples and objects like Parsons' cloth *do* seem possible. I agree that they seem possible under the orthodox view. After all, I can imagine the black and white rectangle in Figure 3, subfigure (a), existing without the two squares. Similarly, I can imagine a cloth like the one Parsons has in mind without any distinct point-sized proper parts. But put the orthodox view out of your mind. Can you really imagine a black and white rectangle which you wouldn't be inclined to describe as "having a black half and a white half"? If we are going to rely on how things seem to be, then I think we should indeed conclude that extended simples with heterogeneous properties are impossible.

Parsons' second case is more difficult since there is more reason to worry about the conclusion. There is, I think, some reason to be suspicious of material objects which have zero volume.[23] Thus, instead of adopting a hard line in this case I will leave things open. What I will say is this. If you think that Parsons' cloth seems to have point-sized proper parts of different colours, then you ought to think it is impossible for it to lack point-sized parts. If, on the other hand, you don't think we can make sense of point-sized objects, then you ought

[22] That such properties are possible is not as obvious as it may seem. The fact that one can imagine a colour gradient like one sees in computer graphics programs, for instance, is misleading since those colour gradients are not truly continuous. Nor is the fact that colour appears to be a quantitative property reason to think that continuous variation in colour across space is possible. Those are not the same thing.

[23] See Arntzenius and Hawthorne (2005) for discussion of this and related issues.

to think that the cloth lacks point-sized parts, and necessarily so. Importantly, this latter stance is not inconsistent with what I have argued in this paper. The key claim, for the purposes of this paper, is not that *all* heterogeneous properties of composite objects are reducible to the homogeneous properties of their proper parts, but that *many* are, including the property *black and white*. In fact, even if *black and white* were the only reducible heterogeneous colour property, my argument would still go through.

Nothing I have said entails that there are no irreducibly heterogeneous properties. Nor does composition as identity entail it. So long as there is good reason to treat some heterogeneous properties differently than others—because, for example, some heterogeneous properties of composites don't seem to reduce to homogeneous properties of their proper parts whereas others do—then we may do so.

There is more to be said on this, but here is not the place. We have said enough to establish that the arguments, as they stand, are no great threat to the view I am defending. Those arguments require that we have more reason to believe that extended black and white simples are metaphysically possible than we have to believe that describing something as "black and white" is a description of its proper parts. That simply does not seem to me to be the case.

CONCLUSION

We have now surveyed what appear to be the best possible accounts of the properties of composite objects. Our discussion suggests that there are three main alternatives: (1) accept composition as identity and the Lewisian account of heterogeneous properties; (2) accept the orthodox view and deny that composite objects have the heterogeneous properties we typically attribute to them;[24] or (3) accept the orthodox view together with anti-reductionism about heterogeneous properties.

All three alternatives are costly, at least from some reasonable viewpoint. We have seen that accepting composition as identity allows for an elegant account of the properties of composites, and one which fits nicely with how we ordinarily think about them. However, many philosophers are reluctant to accept the views about persistence, *de re* modal properties, and quantification that go along with it.[25] To them, the cost of accepting composition as identity seems high.

The second alternative strikes me as completely untenable. Perhaps the simplest argument against it is the best: composite objects *do* have colour properties, as well as all sorts of other non-relational heterogeneous properties. Giving up on this is a step too far.

The third alternative is better, but not by much. Again, the simplest objection may well be the strongest: how could heterogeneous properties like *black and white* not be reducible to properties like *black* and *white*? What's worse, this view seems to commit us to brute similarity relations between heterogeneous properties, as well as irreducible orientation properties. How much nicer it is to stick with the Lewisian view.

I think the immediate consequences of this discussion are clear. When it comes to accounting for the properties of composite objects, composition as identity far outdoes the orthodox view. The broader consequences are less clear. Composition as identity may do better when it comes to accounting for the properties of composite objects, but it also raises new challenges of its own. I hope to have shown that there is value in trying to meet those

[24] If *black and white* is, by definition, a colour property, then, on this view, no composite object is black and white. If *black and white* is not, by definition, a colour property, then, on this view, some composite objects are black and white, but *black and white* is not a colour property. Either way, the view entails that composite objects do not have heterogeneous colour properties. The same goes for other kinds of heterogeneous properties.

[25] See Wallace (2011a, 2011b) for an overview of these issues, and some possible responses.

challenges. In the end, which view comes out on top depends in part on other debates not touched upon in this paper.

As I have said, my money is on composition as identity.

REFERENCES

Arntzenius, Frank, and Hawthorne, John (2005). “Gunk and Continuous Variation”. *The Monist* 88(4), pp. 441–465.

Bøhn, Einar Duenger (2012). “Monism, Emergence, and Plural Logic.” *Erkenntnis,* 76, pp. 211–223.

Botterell, Andrew (2004). “Temporal Parts and Temporary Intrinsics.” *Metaphysica*, 5 (2), pp. 5–24.

Braddon-Mitchell, David and Miller, Kristie (2006). “The Physics of Extended Simples.” *Analysis,* 66, pp. 222–226.

Cameron, Ross (2014). “Parts Generate the Whole, But They Are Not Identical to It.” In Aaron Cotnoir and Donald Baxter (eds.) *Composition as Identity*. Oxford: Oxford University Press.

Cotnoir, Aaron J. (2013). “Composition as General Identity.” In Karen Bennett and Dean W. Zimmerman (eds.) *Oxford Studies in Metaphysics*. Vol. 8. Oxford: Oxford University Press.

Ehring, Doug (1997). “Lewis, Temporary Intrinsics, and Momentary Tropes.” *Analysis*, 57, pp. 254–258.

Gilmore, Cody (2014). “Location and Mereology.” In Edward N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*. Fall, 2014. URL: <http://plato.stanford.edu/entries/location-mereology/.>

Haslanger, Sally (1989). “Endurance and Temporary Intrinsics.” *Analysis,* 49 (3), pp. 119–125.

Jaeger, Andrew J. (2014). “A Tale of Two Parts.” *Res Philosophica*, 91 (3), pp. 477– 484.

Johnston, Mark (1987). “Is There a Problem About Persistence?” *Proceedings of the Aristotelian Society, Supplementary Volumes*, 61, pp. 107–135.

Lewis, David (1983). “Extrinsic Properties.” *Philosophical Studies,* 44, pp. 197– 200.

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

Lombard, Lawrence B. (2003). “The Lowe Road to the Problem of Temporary Intrinsics.” *Philosophical Studies,* 112, pp. 163–185.

Lowe, E. J. (1987). “Lewis on Perdurance versus Endurance.” *Analysis*, 47 (3), pp. 152–154.

Markosian, Ned (1998). “Simples.” *Australasian Journal of Philosophy*, 76, pp. 213– 226.

Markosian, Ned (2004). "Simples, Stuff, and Simple People." *The Monist*, 87, pp. 405–428.

McDaniel, Kris (2003). "Against MaxCon Simples." *Australasian Journal of Philosophy*, 81 (2), pp. 265–275.

McDaniel, Kris (2007a). "Brutal Simples." In Dean W. Zimmerman (ed.) *Oxford Studies in Metaphysics*. Vol. 3. Oxford: Oxford University Press.

McDaniel, Kris (2007b). "Extended Simples." *Philosophical Studies*, 133 (1), pp. 131–141.

McDaniel, Kris (2008). "Against Composition as Identity." *Analysis*, 68 (2), pp. 128–133.

McDaniel, Kris (2009). "Extended Simples and Qualitative Heterogeneity." *The Philosophical Quarterly*, 59 (235), pp. 325–331.

Merricks, Trenton (1999). "Composition as Identity, Mereological Essentialism, and Counterpart Theory." *Australasian Journal of Philosophy*, 77, pp. 192–195.

Parsons, Josh (2004). "Distributional Properties." In Frank Jackson and Graham Priest (eds.) *Lewisian Themes: The Philosophy of David K. Lewis*. Oxford: Oxford University Press.

Sider, Theodore (2007a). "Against Monism." *Analysis*, 67, pp. 1–7.

Sider, Theodore (2007b). "Parthood." *Philosophical Review*, 116, pp. 51–91.

Simons, Peter (2004). "Extended Simples." *The Monist*, 87 (3), pp. 371–385.

Spencer, Joshua (2010). "A Tale of Two Simples." *Philosophical Studies*, 148, pp. 167–181.

Tognazzini, Neal A. (2006). "Simples and the Possibility of Discrete Space." *Australasian Journal of Philosophy*, 84 (1), pp. 117–128.

Van Inwagen, Peter (1990a). "Four-Dimensional Objects." *Noûs*, 24, pp. 245–255.

Van Inwagen, Peter (1990b). *Material Beings*. Ithaca and London: Cornell University Press.

Wallace, Megan (2011a). "Composition as Identity: Part 1." *Philosophy Compass*, 6 (11), p. 804–816.

Wallace, Megan (2011b). "Composition as Identity: Part 2." *Philosophy Compass*, 6 (11), pp. 817–827.

Yi, Byeong-uk (1999). "Is Mereology Ontologically Innocent?" *Philosophical Studies*, 93, pp. 141–160.