

Against the Maximality Principle

C. S. Sutton

Abstract: To hold that only one conscious thing is sitting in your chair, philosophers have appealed to *maximality*: If a property *M* is *maximal*, then anything that has property *M* does not have large proper parts that have property *M*. Philosophers have said that ordinary objects are maximal, including houses, cats, rocks, and have argued by analogy that consciousness is maximal. I argue that the maximality principle mistakenly excludes some members of a kind. Thus, it is not the correct principle to explain why, for example, you are conscious but the proper part that is all-of-you-but-your-arm is not conscious.

1. Introduction

Consider a large part of your body, specifically, the part of your body that is all of you except your left arm, and call that part of your body 'Arm-less'. Arm-less has a brain and nervous system. If the mental supervenes on the physical, and you and Arm-less are the same neurologically, then Arm-less should be conscious. But it seems wrong to say that two conscious beings remember eating your breakfast this morning.

Along comes the *maximality principle* to save the day! The maximality principle tells us that a conscious being cannot have a large proper part that is also a conscious being, and more generally, that ordinary kinds of things, such as houses and cats, cannot have large proper parts that belong to the same kind. Arm-less is not conscious because the property of *being conscious* is had only by the maximal candidate, and Arm-less is not the maximal candidate—you are. Only one conscious being remembers eating your breakfast, despite the abundance of otherwise-suitable supervenience bases for consciousness. The principle is convenient, and it has been gaining traction. The maximality principle has been recently employed, with varying levels of endorsement, by Bailey (forthcoming), Williams 2013, and Walsh 2011. It was earlier endorsed by Burke 1994, Noonan 1999, and Sider 2001 and 2003. Merricks 1998 and Hawley 1998 do not explicitly endorse it but seem to assume it. I will argue that we should reject it as the solution to the problems of the nested consciousness-candidates and nested everyday objects.

Defining terms

The *maximality principle* holds that ordinary kinds of things and consciousness are *maximal*. A *maximal* property *M* is one such that anything that has *M* does not have large proper parts that also have *M*. Some properties are clearly maximal: the property of being the biggest thing in the room is maximal, for example. Something that is the biggest thing in the room could not have a proper part that also has the property of being the biggest thing in the room. Other properties clearly fail to be maximal: the property of being made of matter is not maximal. Something could both be made of matter and have a large proper part that is also made of matter.

The Maximality Principle's Promised Work

The maximality principle is used to secure in our ontology the correct number of ordinary objects and conscious beings. It tells us that a house does not have large proper parts that are also houses. Consider, for example, the large proper part of a house that is all of the house except the front shutters. The maximality principle secures the result that this near-house is not a house: there is only one house on the lot. The same goes for David Wiggins's Tib, which is the part of the cat Tibbles that is all but the cat's tail. If *being a cat* is maximal, then maximality promises that Tib is not a cat. Both Tibbles and Tib might be ripping up the sofa, but only one is a cat, namely, Tibbles.

In the philosophical literature, the maximality principle has been motivated by appeal to our intuitions about the maximality of kind properties like *being a house* and *being a cat* and is often applied to the property of *being consciousness* or to *being a conscious being*. The maximality principle might well be *de facto* true in most cases of kinds. It would be inefficient and impractical to build telephones using other telephones as parts, and I suspect that the same goes for most kinds. But there are some kinds that cause problems if we apply the maximality principle to them and also cause problems if we claim that they are exceptions to the maximality principle.

I will argue that the maximality principle does not consistently do the promised work of securing the desired answer about how many kind tokens there are. I present three cases of ordinary objects in which the maximality principle gets the wrong answer. Because the maximality of consciousness is motivated by analogy to the maximality of ordinary kinds of

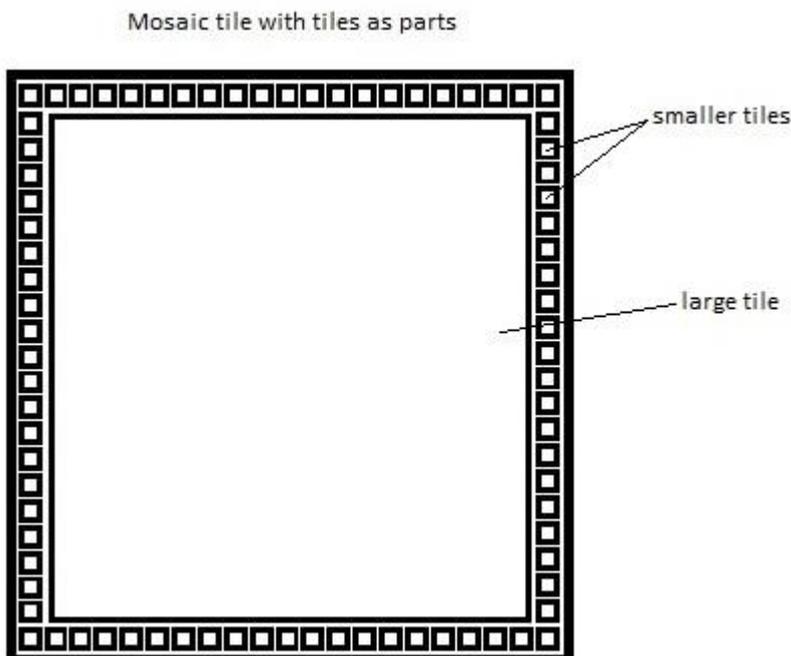
things, the wrong answer about everyday kinds should undercut our confidence that the maximality principle is the correct explanation for why Arm-less is not a conscious being.

Note that the maximality principle rules out only *large* proper parts from having the same maximal property as the whole (Sider 2001, 357). For example, a tile manufacturer might offer tile squares that are mosaics made with smaller tiles. Both the mosaic tile and the smaller tiles in the mosaic have the property of *being a tile*. You, as an organism, have smaller organisms as parts of your digestive tract. The organisms in your digestive tract are small parts of you, so the maximality principle does not rule them out from organism-hood. The proper part of you that we called Arm-less, on the other hand, is a large part, so the maximality principle rules that Arm-less is not an organism. Even though the largeness constraint seems helpful in these cases, we see that in the cases below, it does not do the desired work—it is overzealous and undercounts the number of kind tokens.

2. Problem Cases for the Maximality Principle

The Tile Case

Consider a mosaic tile in which the mosaic pattern has a large tile piece surrounded by smaller tiles.



In such a case, the maximality principle would mistakenly have it that either (1) the large part of the mosaic is not a tile even though the smaller pieces are tiles or (2) the large part *is* a tile but the object that it is a part of, the mosaic tile, is *not* a tile. The maximality principle gets this case wrong: the mosaic tile and its large tile part are both clearly tiles.

The Coat Case

If you are looking for a heavy-duty coat to withstand winter storms like we have seen this year, you might consider a coat which has two coats that zip together to make a heavy-duty coat. On fall days, you might use only the lighter coat, which makes the inner layer. Other days you might use the heavier coat-part by itself, with the inner lighter-weight coat removed. You would use the whole coat, with the lightweight and the heavy coat zipped into one, if you went skiing. Imagine such a coat in which the lighter layer is very light; it makes up five percent of the volume of the entire put-together coat. The heavier layer makes up 95% of the coat. (If you think that 95% is not high enough to count as a large part, then imagine a coat fit for Antarctica, where the inner layer is the lightest type of fall coat, and the outer layer makes up an even higher percentage of the put-together coat.) If *being a coat* is a maximal property, then either (1) the heavier layer is not a coat, or (2) the put-together layers are not a coat. But both the large outer layer coat and the whole zipped-up coat look like and function as coats. The maximality principle gets the coat case wrong.

The File Case

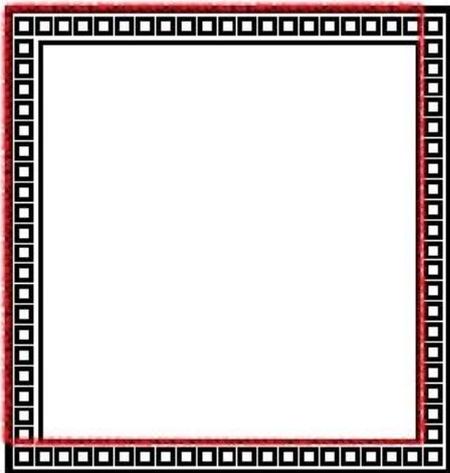
In my desk drawer, I have files, and those files contain smaller files. I have a research file, and in that file, there is a file for each paper that I am working on. The files for each paper are parts of the research file. We can imagine a file for conference organizing, and imagine that I have organized exactly one conference, the 2014 Logic Lovers Conference, so the conference file contains only one file (the 2014 Logic Lovers Conference file), and that file is rather large. The conference file might in the future grow to include files about other conferences, though the 2014 Logic Lovers Conference file would not grow with such an addition, which suggests that the conference file cannot be reduced to the 2014 Logic Lovers Conference file. If the maximality principle were true, then a file could not have another file as a large part, so either (1) the 2014 Logic Lovers Conference file is not a file or (2) the conference file is not a file. But both the files are clear cases of files; neither is a

fringe or borderline case of a file. Once again, the maximality principle gets the case wrong. It miscounts the number of files as one rather than two.

Objection: Tilehood, coathood, and filehood are not maximal properties

We might be tempted to reply that the properties of *being a tile*, *being a coat*, and *being a file* are not maximal. For the maximality principle makes no claim that *all* properties are maximal, only that if they are, they cannot be had by both things and their large proper parts. If *being a tile* is *not* maximal, then the large tile-part is rightly a tile. But so would be 'tile-minus', the proper part of the tile shown below, whose perimeters are indicated by the fuzzy red line.

Tile-minus, a Proper Part of the Mosaic Tile



Tile-minus is just sort of thing the maximality principle is designed to disqualify from our ontology. And tile-minus is just one example. We could redraw the red line in various other ways to make large proper parts that, like tile-minus, are not intuitively tiles. If we abandon the maximality of tilehood, then we will have a swarm of overlapping and nested tiles like tile-minus. Consider also Coat-minus, which is composed of the zipped-together light and heavy layers, minus the sewn-on pockets, both exterior and interior pockets. If coathood is not maximal, then Coat-minus would qualify as a coat. But coat-minus is just the kind of thing that the maximality principle is designed to qualify from our ontology. Coat-minus has cousins, such as all of the zipped-together coats except the collars. If we take coathood as

an exception to the maximality of kinds, then we will have a multitude of nested and overlapping coat candidates that qualify as coats.

Remember that the maximality principle is endorsed in order to root out from our ontology nested groups of cats, houses, and the like. But if we use exceptions to the maximality principle in order to get around cases like the tile, coat, and file cases, then our ontology will have pockets of the unsightly ontological clutter that the maximality principle is designed to avoid. So denying that the maximality principle applies in the problem cases is not the answer. Furthermore, if we allowed nested groups of tiles and coats and files, then why not allow parallel things like nested groups of houses and cats? The maximality principle is invoked to avoid such nested groups, but because it does not do so systematically, we should keep looking for a more accurate principle that does (or else bite the bullet and allow that there are such nested groups).

3. The Conceptual Maximality Principle

Let us consider a version of the maximality principle that accepts nested clutter, as long as there is *some* sense in which there is only one token *x*. The *conceptual maximality principle* holds that kind predicates and consciousness are maximal *in virtue of linguistic or conceptual truth*. For example, *being a conscious being* is maximal because our conceptual scheme has it that there is only one conscious being eating your breakfast. I will argue that the conceptual maximality principle relies on a mistaken account of our conceptual scheme, but first, let us look at what the conceptual maximality principle says and what motivates it.

Theodore Sider has proposed that proper parts of conscious things, such as Arm-less, are not conscious, but they are ‘psuedo-conscious’. Sider describes this pseudo-consciousness as that “which a thing has in virtue of having all that is required, intrinsically, for consciousness. [Pseudo-consciousness] is consciousness stripped of any maximality requirement” (2003, 147).¹ In contrast with pseudo-consciousness, regular consciousness satisfies our conceptual framework that consciousness is maximal. If we apply the conceptual maximality principle to consciousness candidates, we see that you are conscious but Arm-less is pseudo-conscious. Arm-less is joined by an army of other nested candidates

¹ Sider uses the term ‘consciousness*’ in his published work, rather than ‘pseudo-consciousness’. He used ‘psuedo-consciousness’ in a manuscript, and I will follow Sider’s original terminology of ‘pseudo-consciousness’ because it is easier to discriminate visually between the terms ‘conscious’ and ‘pseudo-conscious’.

that fall short of *being conscious* but have the property of *being pseudo-conscious*. Pseudo-consciousness is not maximal, but that does not run afoul of our concept or linguistic convention of pseudo-consciousness because *there is no* concept or linguistic convention about pseudo-consciousness. *Mutatis mutandi* for pseudo-houses, pseudo-cats, and so on.

Sider sums up the view thus: “it's a conceptual truth that something counts as conscious (or, as a conscious being) only if it's the ‘largest’ conscious thing in the vicinity. The concept of *being conscious* is a maximal concept. If our evidence that there are not thousands of conscious beings in Martha's immediate vicinity is not conceptual in this way, then it's difficult to see what that evidence could be” (Sider 2003, 144–145).

The conceptual maximality principle has two motivations. The first, shared by the non-conceptual version, is that philosophers want to get counting and reference correct. Assume that only one person, you, ate your breakfast this morning. The maximality principle is supposed to give us a principled reason why there was only one person and why you, and not Arm-less or one of its ilk, is that person.

The second motivation for the conceptual maximality principle is to preserve the supervenience of objects and their properties on the parts of those objects and the relations of those parts (including intrinsic and extrinsic relations). If mental states supervene on physical states, then if you love Dámaso Alonso's poem “The Star Counters,” then we would expect Arm-less to love the poem, too. After all, Arm-less has the same brain that you do, and it has experienced the same events that you have. The left arm does not seem to be involved in one's attitudes toward poetry. The conceptual maximality principle allows that both you and Arm-less have the mental state of love for the poem, but that only you are ‘conscious’ and a ‘person’. You might wonder how something that is not conscious can appreciate a poem, but remember that pseudo-consciousness is just like consciousness except insofar as consciousness is maximal.

The conceptual maximality principle succeeds in the second motivation, preserving supervenience as described above, but it fails to satisfy the first motivation of correctly counting how many tokens of a kind there are. This failure is illustrated by the same cases that cause a problem for the non-conceptual version of the principle: the tile case, the coat case, and the file case. The conceptual maximality principle tells us that there should only be one large tile (the mosaic tile) with many smaller tiles as parts. It excludes from the count

the largest tile used to make the mosaic. Likewise, the conceptual maximality principle omits the large-coat part and the large-file part from the count of coats and tiles.

According to the conceptual version, it is part of our concept of tilehood that the large-tile part is not a tile, the large-coat part is not a coat, and the large file-part is not a file. Sider takes linguistic intuition as evidence of our conceptual content. But the evidence from linguistic intuition is that people are happy to call these large parts members of the respective kind of thing. “Please pass me the Logic Lovers Conference file” does not illicit the response, “That is not a file—it is only part of a file.” If you point to the largest tile in the mosaic and say, “I like the texture of this tile the best,” no one will be confused and think that you have misused the word ‘tile’. This is evidence that our concept of ‘tile’ is not what the maximality principle claims it to be.

Conclusion

The maximality principle does not do the promised work of securing all of the commonsense cases of kind membership. Even if the maximality principle *de facto* applies in most cases, it is not systematically accurate, and using the maximality principle to solve our counting problems will distract us from finding the true reason that you count as a conscious being but Arm-less does not.

References

- Bailey, Andrew (2014) You Needn't Be Simple. *Philosophical Papers* 43: 145–160.
- Burke, Michael (1994) Dion and Theon: An Essentialist Solution to an Ancient Puzzle. *The J. of Philosophy* 91, no. 3: 129–139.
- Hawley, Katherine (1998) Merricks on Whether Being Conscious is Intrinsic. *Mind* 107: 841–843.
- Merricks, Trenton (1998) Against the Doctrine of Microphysical Supervenience. *Mind* 107: 59–71.
- Merricks, Trenton (2001) *Objects and Persons*. New York: Oxford University Press.
- Merricks, Trenton (2003) Maximality and Consciousness. *Philosophy and Phenomenological Research* 66: 150–158.
- Noonan, Harold (1999) Identity, Constitution, and Microphysical Supervenience. *Proc. of the Aristototelean Society, New Series, Vol. 99*: 273–288.
- Sider, Theodore (2001) Maximality and Intrinsic Properties. *Philosophy and Phenomenological Research* 63: 357–364.
- Sider, Theodore (2003) Maximality and Microphysical Supervenience. *Philosophy and Phenomenological Research* 66: 139–149.
- Van Inwagen, Peter (1981) The Doctrine of Arbitrary Undetached Parts. *Pacific Philosophical Q.* 62: 123–137.
- Van Inwagen, Peter (1990) *Material Beings*. Ithaca, New York: Cornell University Press.
- Walsh, Sean Drysdale (2011) Maximality, Duplication, and Intrinsic Value. *Ratio* 25: 311–325.
- Williams, J. Robert G. (2013) Part-Intrinsicity. *Noûs* 47, no. 3: 431–452.