

Chapter 5

Advantages of CI

1. Introduction

By now, I have hopefully convinced you that despite the many objections against it, CI is a coherent view worth defending. In Chapter 2, I have shown how a CI theorist can defend herself against Van Inwagen's Counting objection, using Plural Counting. In Chapter 3, I demonstrated how a CI theorist could defend herself other objections as well, such as those involving the Indiscernibility of Identicals, those that appeal to the Principle of Ontological Parsimony, etc., by invoking both Plural Counting and a robust plural language. And in the preceding chapter, Chapter 4, I showed how a CI theorist could defend herself against the Modal Objection, by embracing a lump theory of objects.¹ All of this should at least show that CI is metaphysical view to be taken seriously, and one that can't be dismissed out of hand merely because it seems incoherent on its face.

However, I would like to do more than merely make room for CI in logical space. I would like to illustrate, by way of a few (more) puzzles, how advantageous such a view can be. In the previous chapter, we invoked the strategy of measuring the success of a view by its utility. If (e.g.) the lump theory of objects can solve a myriad of constitution puzzles, then this is some reason to think that the theory is true. I would like to employ that same strategy here. Only, I will no longer be

¹ I also (quickly) showed two alternative ways a CI theorist could defend herself against the Modal Objection *without* embracing my lump theory of objects—*viz.*, by either embracing a counterpart theory, and admitting high flexibility in *de re* predication, or by embracing a theory of objects (suggested by L. A. Paul) where objects are mereological sums of properties.

concerned merely with puzzles of composition and constitution; CI has a much broader appeal than this. This is primarily due to the fact that once CI is accepted, then one can appeal to mereological sums with impunity, because (e.g.) mereological sums are no longer an ontological burden. And this should be as expected. After all, we began this thesis with the worry about whether mereological sums were indeed ontologically innocent or not. I have defended CI so vigorously because it easily delivers the verdict that, yes, mereological sums *are* ontologically innocent: if you are already committed to a bunch of things, then you get the mereological sum of those things *literally* for free, because the sum is simply *identical* to those things. Of course, you might now be thinking, “OK. Fine. I see that CI is coherent, and that none of the traditional arguments against are effective. But so what? Even if CI is true, what good does having (ontologically innocent) mereological sums in our ontology do? ”

What good it will do is exactly the point of the present chapter. I want to show, by way of a few puzzles, how accepting CI and embracing mereological sums will deliver wide-ranging application to many different areas in philosophy. This won't be a comprehensive list of *all* of the philosophical areas in which CI will be prove to be fruitful. But I will canvass just a few of them—e.g., areas in causation, prevention, and perception—and demonstrate how problems in these areas could be benefitted by appealing to CI, in virtue of the fact that CI delivers mereological sums burden-free. I will also hint at CI's application to other areas (e.g., moral responsibility, philosophy of mind, etc.). The general goal of this chapter, then, is to answer the “so what?” question. I will provide a template of problems that litter the philosophical

landscape, and show how CI, and an appeal to mereological sums, offers a solution to them.

My strategy will be to first discuss two classic cases of causation: collective causation and overdetermination. Then I will lay out four philosophical puzzles and show how all of them are similar in structure, and can be taxonomized *in part* as a case of collective causation, and *in part* as a case of overdetermination; the cases of collective causation and overdetermination will be used to frame the four puzzles I introduce below. Then I want to show how adopting CI (and mereological sums) provides a unified, elegant solution to all of them. In doing so, it should then be easy to show that CI is not only merely available as a coherent metaphysical thesis, but that its utility and wide-ranging applications to various areas in philosophy should give us great reason to think that it is true.

2. Some Taxonomy

There are two kinds of causation I want to present: *collective causation* and *overdetermination*.

To illustrate a classic case of collective causation, imagine that two men, ONE and TWO, are given the task to set off a bomb. The bomb is designed so that two men must turn two separate keys at the same time. If ONE turns the key and TWO does not, the bomb will not go off. If TWO turns the key and ONE does not, the bomb will not go off. (And let us assume that no one except ONE and TWO have access to the keys.) So, ONE and TWO each reach for their respective keys, and turn them simultaneously. Predictably, the bomb goes off.

Now, who caused the bomb to go off: ONE or TWO? In a way, this is a misguided question, if the disjunction is read exclusively. Intuitively, ONE wasn't *the* cause of the bomb going off, since if TWO would not have turned the key, ONE's actions would have been causally ineffective (as far as a detonation of the bomb is concerned). But TWO wasn't *the* cause of the bomb going off either, since if ONE would not have turned the key, the bomb would not have gone off. So neither ONE nor TWO, taken *individually*, were sufficient for the bomb exploding. But our question—*who caused the bomb to go off?*—seems to imply that there is *one* person who caused the bomb to go off. But surely someone (or something) caused the bomb to go off! It didn't just explode causelessly! But then what?

We don't usually take the above case to be problematic. This is because we understand that individuals can act collectively to cause something to happen. If Rod and Todd lift the coffin collectively, such that neither Rod nor Todd could lift the coffin by themselves, then we understand that it is Rod and Todd together who *caused* the coffin to be lifted. What we have said in previous chapters about objects—plural—doing certain things collectively applies to causation as well. And, indeed, this notion is plenty familiar. I mention the case here so that we can catalogue it in comparison with the puzzles I will be introducing below. Before we evaluate the case of collective causation in comparison with the puzzle cases, however, let us take a look at a classic case of overdetermination as well.

Imagine that Man₁ and Man₂ have the unfortunate task of killing Man₃. Each is given a loaded gun, and each has excellent aim. Both shoot, and each shot is fatal: If Man₁ had not shot his gun, but Man₂ had, Man₃ would have died. If Man₂ had not

shot his gun, but Man₁ had, Man₃ still would have died. A familiar way of describing the case is that each of Man₁ and Man₂ was sufficient for killing Man₃, but neither (taken individually) was necessary.

Overdetermination cases such as these² are often used as counterexamples to counterfactual theories of causation.³ In the case of Man₁ and Man₂, there will be difficulty saying that either one killed Man₃. Man₃ would have died anyway if Man₁ would have shot him, but Man₂ did not. So it doesn't seem that Man₂ *did* anything (causally relevant, anyway). Likewise, Man₃ would have died anyway if Man₂ had shot him but Man₁ had not, so it doesn't seem that Man₁ did anything either. Man₁ and Man₂, for parallel lines of reasoning, then, both seem causally *irrelevant*. So the worry here is not so much that we would deny the existence of either Man₁ or Man₂ (or the existence of the bullets fired from Man₁ or Man₂'s gun), as we might in other cases of overdetermination (e.g., mind/body dualism, or composite objects—see below). Rather, the worry is that once we have admitted that one of the shooters is sufficient for the death of Man₃, then the other shooter becomes seemingly causally irrelevant, and hence, not a cause at all. Since this counterfactual reasoning is

² I do not mean to suggest that this case is the only kind of overdetermination, and indeed, there may be important differences between various different kinds of overdetermination. See Funkhouser (2002). However, I hope that for my purposes here, we can ignore the subtle distinctions of the different kinds of overdetermination there may be, and simply consider overdetermination in general.

³ Overdetermination is also seen as a counterexample to nomic subsumption theories of causation, which claim that a property (usually a mental property) can be causally relevant if it appears under a law (e.g., mental-physical laws). Cases of overdetermination are seen as counterexamples to it because (e.g.) we might have a mental event and a physical event that both are sufficient for causing another physical (or mental) event, and this would undermine the intuition that (e.g.) the mental event is causally relevant. In the interest of time, I am going to omit discussion of this issue, but it may become obvious in the course of the discussion, what I *would* say about such a case. I am going to be arguing that overdetermination shouldn't trouble us in general, and so I do not think it should be a threat to nomic subsumption theories of causation either. Of course, there may be other reasons not to endorse such a causal theory, but these won't be investigated here.

symmetric, we get the odd result that neither Man_1 nor Man_2 killed Man_3 . This is why overdetermination seems problematic in the case of the two shooters: having two sufficient causes seems to lead to the absurd conclusion that each was causally irrelevant, or not a cause at all.⁴

Overdetermination is also seen as problematic in general, and a seeming violation of simplicity principles in theory building. After all, if one cause is sufficient for explaining an event, then it would be ontologically excessive to posit another, distinct cause in addition. For example, in philosophy of mind, many have thought that the Exclusion Problem shows that there cannot be mental causes in addition to physical causes, if the physical causes are sufficient for bringing about a particular event.⁵ Mental causes are given up because to embrace them would unnecessarily proliferate instances of overdetermination, which is theoretically inelegant.⁶

As another example, Trenton Merricks argues for an Eliminativist view of ordinary objects (e.g., tables and chairs) in order to avoid rampant overdetermination.⁷ Merricks claims that if composite wholes existed, they would overdetermine events which are sufficiently caused by the composite's parts—e.g., the parts of a baseball cause a window to shatter, and are sufficient for the shattering, so it would be causally redundant to claim that the (whole) baseball causes the shattering as well. Since the baseball example generalizes to all cases of

⁴ Of course, this assumes a counterfactual account of causation, which may be given up if overdetermination cases such as the one involving the two shooters want to be embraced as coherent possibilities.

⁵ See: Kim (1993c), Schiffer (1987), Merricks (2001), etc.

⁶ Merricks says, in response to overdetermination worries of the mental and the physical, "The redundancy *all by itself* is reason to resist...substance dualism." Merricks (2001: 67).

⁷ Merricks (2001).

part/whole causation, we should deny that there are any wholes. Such rampant causal redundancy, or overdetermination, in other words, should be avoided at all costs, and we would do better to deny that there are (e.g.) baseballs and running shoes than that overdetermination is ubiquitous. Merricks claims, "...we always have a reason to resist systematic causal overdetermination, along with any view that implies it."⁸ So clearly, overdetermination is seen as problematic in general, and something to be avoided as often possible.

My purpose for bringing up overdetermination here is two-fold: I want to compare some of the puzzle cases with them, and I want to eventually argue that these may not be effective counterexamples to counterfactual analyses of causation after all, nor—contra Merricks—should rampant overdetermination necessarily be avoided in general. If we adopt CI and we admit that mereological sums can be causes, then our aversion to overdetermination might be mitigated. I will get to these two points further down below. Let us keep collective causation and overdetermination in mind, however, and take a look at some philosophical puzzles.

3. Four Puzzles

There are four puzzles I would like to present: (i) Shadow, (ii) Eclipse, (iii) Prevention, and (iv) Perception. I will present these below, and suggest how they might be categorized given our two causation cases above, and how they can be solved by appeal to mereological sums. After I have discussed these, I will gesture at a few other, similarly structured puzzles—cases involving moral responsibility,

⁸ Merricks (2001: 67)

Frankfurt Cases, the Exclusion Problem in philosophy of mind, etc.—and show how these, too, might be solved by appeal to mereological sums.

i. *Shadow*⁹

Imagine that there is a light source aimed at an opaque disk, O_1 . Several feet behind O_1 is another opaque disk, O_2 . O_2 is in just the right position, and just the right size, such that if O_1 were not in front of it, and if the light source had hit O_2 , then a shadow would be cast on the ground that is the exact size and shape as the shadow that is on the ground when O_1 is in place. So, if O_1 stays where it is, and O_2 is removed, the shadow is a certain shape and size, s . If O_1 is removed, but O_2 is in place, then there is a shadow of the certain size and shape, s . But let's return to the case where both O_1 and O_2 are in place, the latter several feet behind the former, as illustrated by Figure 1:




Figure 1

Now the question is: which object, O_1 or O_2 is casting the shadow? Intuitively, it is not O_1 , because O_1 's shadow is seemingly blocked by O_2 . But it is not O_2 's shadow either, since O_2 is blocked from light; intuitively, an object must be in the

⁹ This example modified from Sorensen (2006).

way of a light source, not in complete darkness itself, in order for it to cast any shadow at all. So it seems that neither O_1 nor O_2 are casting the shadow. But then *what is?* Clearly there is a shadow there!

We can contrast our intuitions about a situation involving two objects, as illustrated by Figure 1, with the usual, accepted case where there is only one object involved. Imagine that due to some mysterious physical event, O_1 and O_2 begin to move closer and closer to each other, and then begin to fuse into a single opaque object, where O_1 is the front half of the object, which faces the light source, and O_2 is the backside of the object which is closest to the cast shadow. But then this is just the ordinary case where one opaque object is causing a shadow! Ordinarily, when we have an instance of something causing a shadow—a building, a person, a wall, etc.—the object is almost always thick enough to have a front half and a back half. But just because O_1 and O_2 have fused together, this should intuitively make no difference as far as the difficulty of answering the question: *which object is causing the shadow?* If we don't have an answer in the case of two opaque objects, some distance apart from each other, it seems we do not have an answer in the ordinary case either. In other words, admitting that there is a problem in the case of two objects infects the ordinary case of one object. And so it seems that we lack a satisfactory explanation of how it is that a single object casts a shadow, if we don't have an answer to what object(s?) cast a shadow in Figure 1, contrary to our ordinary intuitions.

Clearly, what I want to say in the Shadow case, given my thorough defense of CI, is that it is mistaken of us to assume that the shadow-caster in Figure 1 must be

either O_1 or O_2 , *individually*. Rather, I will claim that *both* O_1 and O_2 can cast a shadow, *collectively*, without either one of them casting a shadow on their own. Indeed, in the ordinary case of just one object casting a shadow, we do not hesitate to say that the whole object casts the shadow, even if we readily admit that the object in question has a front half and a back half. Modus tollens-ing the above point: if we have no trouble saying that a whole (made up of a front and back half) can cast a shadow—if we have no trouble saying that the front and back half of an object can cast a shadow *collectively*—then we should have no trouble saying that O_1 and O_2 in Figure 1 can cast a shadow *collectively*.

Pushing the point further, we can reverse our case of gradual fusion above, and consider instead a case of gradual fission: take the ordinary case of one object casting a shadow, and divide the object into a front and back half.¹⁰ Then we could slowly, bit by bit, remove the front half away from the back half. It would seem arbitrary to claim that one object could cast a shadow one moment, but that when the front half and the back half of this object are separated by a minute distance, it doesn't. Why wouldn't it, given that there *is* a shadow there, whether the halves are attached or not, and one that is seemingly indiscernible from the shadow that was there before the separating of the front half from the back half? Yet if the front half and the back half form a shadow in the ordinary case (i.e., when they're touching), and a minute separation of the front half and the back half do not make a difference as far as the shadow is concerned (i.e., the halves are still casting a shadow even if they are minutely separated), then a step-wise argument will show us that the halves

¹⁰ Let's assume that these halves are symmetrical.

significantly separated will still (collectively) cast a shadow.¹¹ Yet the case where we have two partitioned halves, significantly separated, just *is* our original Shadow puzzle, modulo a particular history of the two shadow-casters (i.e., they used to be two halves of a conjoined whole).

Now perhaps one might claim that there is an important difference between the front and back half of a (whole) object collectively casting a shadow, and two unconnected objects such as O_1 and O_2 collectively casting a shadow (or two separated, unconnected halves collectively casting a shadow). Perhaps, one might argue, the difference lies in the fact that front and back halves of (whole) objects are attached or connected, whereas objects such as O_1 and O_2 are clearly *unattached* and *unconnected*. And it is because front and back halves (of whole objects) are attached and connected that we have no problem saying that they can collectively cast a shadow, but that two (distinct, separated) objects cannot.

This is incorrect for two reasons. First, recall our case of collective causation: we can allow that two men may cause a bomb to go off, even when neither of the men are connected or attached in any way. Casting a shadow is a causal event: it is causing an absence of light. If we can allow collective causation when the objects involved are unattached or unconnected (such as the two men) then we should allow collective shadow-casting when the objects involved are unattached or unconnected (such as O_1 and O_2). If we do not, then we either need to say why casting a shadow is significantly different than (e.g.) causing a bomb to go off, or we need to say why

¹¹ By 'significantly separated' I mean enough to generate the kind of puzzle with which we began this section. I suspect there is a limit to the distance two (or more) objects could be separated and still cast a shadow—i.e., if the two objects are too far apart from each other, or too far away from the light source, then there won't be any shadow at all. I think we can ignore these details without detriment to the point being made here.

attachment or connection is required in one case but not the other. Either option, I'm afraid, would run the risk of being ad hoc.

Second, however, is the fact that we allow objects to act collectively in all sorts of ways: Rod and Todd can meet for lunch, carry a coffin, enter a three-legged race, hold hands, sing harmony, etc. It would be strange to admit that two unconnected and unattached objects¹² could do so many things collectively *except* cast a shadow. What is so special about casting a shadow that unattached and unconnected items couldn't do it together? If O_1 and O_2 can move in synch, can be symmetrical, can take up a certain region of space, can do-si-do (if they were so inclined)—all *collectively*—then they should certainly be able to cast a shadow (collectively).

Now, you might well agree that we allow objects to engage in activities collectively. You might agree that Rod and Todd can meet for lunch, carry a coffin, etc. But this is different than O_1 and O_2 casting a shadow (as illustrated by Figure 1), you might argue, because *each* of O_1 and O_2 could have cast the shadow all by itself, without the help of the other. Not so with Rod and Todd and all of the activities they can collectively engage in—Rod can't meet for lunch by himself, or carry a coffin, or enter a three-legged race by himself; he needs Todd's help. Similarly, you might point out, O_1 and O_2 might be able to collectively cast a shadow only if in so

¹² Granted, Rod and Todd might have to be connected or attached in some sense in order to (e.g.) run a three-legged race or hold hands, etc., but I take it that this is not the sort of connectivity or attachment someone pursuing this line of argument had in mind. That is, I take it that if O_1 and O_2 were 'touching' at the very edge (suppose they are spheres that are resting side by side for a moment, for example), but still casting a shadow, the imagined objector above would still find the case problematic.

doing they are casting a shadow such that the shadow could not be cast without *both* of them.

So, for example, if O_1 was above O_2 , or slightly to the right or left of it, then the shadow might look like a filled-in figure eight (if O_1 and O_2 were arranged in just the right way). And *this* shadow, you might claim, is one that O_1 and O_2 *can* cast collectively, like Rod and Todd might tango collectively. For if either O_1 or O_2 were to be removed, then the figure-eight shadow would no longer be cast (since one half of the 'eight' would disappear along with the removal of either O_1 or O_2).

But if you grant me that O_1 and O_2 can cast a shadow collectively in the case where they are arranged side by side, then we can make a step-wise argument to show that they are casting a shadow collectively in a case such as that represented by Figure 1. Simply imagine that O_1 and O_2 are side by side and are (collectively) casting a figure-eight shadow. Now slowly move O_1 behind O_2 . At each moment, the shadow will change shape. But so long as O_1 is not *completely* occluded by O_2 , we should still grant that O_1 and O_2 are collectively casting a shadow (since *whatever* the shape of the shadow that is cast by O_1 and O_2 , such a shadow would disappear or distort if we removed either O_1 or O_2). But this means that even if O_1 and O_2 were not completely occluded, but overlapped *nearly* completely, with perhaps a small sliver of non-overlap, one would (presumably) still claim that they collectively cast a shadow. Then, the next moment, when O_1 is completely occluded by O_2 , we have a situation such as the one represented by Figure 1. But then suddenly O_1 and O_2 are *not* collectively casting a shadow!? What's more, we could imagine that O_1 and O_2 continue along their respective paths, and O_1 now slowly moves away from O_2 on

the opposite side. So now O_1 and O_2 suddenly *are* collectively casting a shadow again?! Strange indeed that objects could (by hypothesis) collectively engage in an activity one moment, not collectively engage in it the next, and then collectively engage in the activity again, when the objects in question aren't doing anything different, except having been moved just a tiny bit so as to be in perfect alignment (and then not) relative to the light source.

Thus, on pain of having to draw such an improbably harsh distinction, between a collective activity and a non-collective activity, we should grant that O_1 and O_2 are (at least) collectively casting a shadow. There might be some differences between (e.g.) two objects dancing a tango and two objects casting a shadow as is represented in Figure 1. (And, indeed, I will discuss these differences more fully below.) But the point is that we can grant that two objects can engage in an activity collectively, whether it be tangoing, carrying a coffin, or casting a shadow, etc. The objection that began this discussion, recall, was concerned with whether the front and back halves of objects are different than two separated objects (such as O_1 and O_2), because the former are attached, while the latter are not. Yet we can now see that attachment or connectivity is beside the point. If we can grant that two unconnected, separated objects can engage collectively in an activity (e.g., tangoing or casting a shadow), then someone will not be able to plausibly maintain that in the ordinary case of one object casting a shadow there is no puzzle (because the back and front are connected), yet in the case of two objects, such as the one depicted in Figure 1, there *is* a puzzle (because the objects involved are not connected). Either

the ordinary case of one object and the case with two objects are both problematic, or else they are both unproblematic.

Clearly, I maintain that neither case is problematic. And we can see this by first granting that, in Figure 1, O_1 and O_2 cast the relevant shadow collectively. Given CI, however, the mereological sum of O_1 and O_2 just *is* (hybrid identical to) O_1 and O_2 taken collectively. And so if O_1 and O_2 engage in some activity collectively, then the mereological sum of O_1 and O_2 engage in this activity. So it is the mereological sum of O_1 and O_2 that casts the shadow in our puzzle, Shadow.

Note at this point that Shadow can not only be solved by an appeal to mereological sums, but that it could also be seen as a direct argument for the *existence* of mereological sums.¹³ The argument might run as follows:

1. There is a shadow in Figure 1.
2. If there is a shadow, then something must cast the shadow.
3. O_1 does not cast the shadow (because O_1 's shadow is cast on O_2).
4. O_2 does not cast the shadow (because there is no light hitting O_2).
5. If something must cast the shadow, and if neither O_1 nor O_2 do, then the mereological sum of O_1 and O_2 must cast the shadow.
6. By 1 and 2, something must cast the shadow.
7. So, by 5 and 6, the mereological sum of O_1 and O_2 must cast the shadow.
8. If something casts a shadow, then this something exists.
9. So, mereological sums exist.

Premise 1-4, and 6-8, are all fairly intuitive.¹⁴ Premise 5, I take it, is the one that is the most contentious. But the argument for Premise 5 is one of elimination: mereological sums are simply the only viable candidate for shadow-casting. Let us

¹³ Thanks to Bill Lycan for discussion here. Also, notice that this isn't my preferred argument for the existence of mereological sums, but is one I am obviously sympathetic to given CI.

¹⁴ Of course, this does not mean that they are completely uncontroversial. Someone could reject Premise 2, for example, and argue for some sort of eliminativist view about (some) shadow-casters. If someone thought that sometimes nothing causes shadows, for example, then there might be reason to reject Premise 2. But I take it that this would be an unusual view, and so I will ignore it here.

imagine that O_1 and O_2 are the only objects in the vicinity (besides the light source, and whatever it is that the cast shadow is cast upon). Then what else besides the mereological sum of O_1 and O_2 *could* cast the shadow, if neither O_1 nor O_2 cast it?

One might be tempted to claim that the *set* of O_1 and O_2 cast the shadow, for example. But sets are abstract entities. As such, they do not interact, causally or otherwise, with the non-abstract, material world. They certainly do not cast shadows.

One might be tempted to think that if O_1 doesn't cast the shadow and O_2 doesn't cast the shadow, then perhaps O_1 and O_2 taken together or collectively cast the shadow, yet without this entailing that the mereological sum of O_1 and O_2 cast the shadow. Someone who does not embrace CI as I've defended her in this thesis, for example, may want to pursue this line. Such a person would grant the lessons learned from Chapter 2 about plural referring terms, and would claim that some objects—plural—may engage (collectively) in an activity (such as shadow-casting) without this activity distributing down to the objects individually.

Fair enough. But if someone is willing to grant me this much, then I will offer all of the arguments for CI in the present thesis to support the claim that if O_1 and O_2 collectively do something—e.g., cast a shadow—then their mereological sum does it, too, since the mereological sum of O_1 and O_2 just *is* O_1 and O_2 taken together or taken collectively. Anything O_1 and O_2 do together, their mereological sum does, given CI. So, if desired, we can break Premise 5 into two parts, 5a and 5b:

- 5a. If something must cast the shadow, and if neither O_1 nor O_2 do, then O_1 and O_2 must cast the shadow collectively.

- 5b. If O_1 and O_2 cast the shadow collectively, then the mereological sum of O_1 and O_2 casts the shadow.

Premise 5a is supported by the reasons outlined above. Premise 5b is supported by CI, which is argued for throughout this entire thesis. So if someone objects to the above argument on the grounds that he grants that objects can engage in activities collectively (e.g. cast shadows), but disagrees that such collective activity entails that a mereological sum engages in said activity, then the above argument (on its own) will not be sufficient to prove the existence of mereological sums. Our debate will then shift to arguments for 5b. But since one way to argue for 5b is simply to endorse CI, then I offer all of the arguments in this present thesis to oblige such an objector.

Another way to resist Premise 5 is as follows: even if one grants that the only viable candidate for shadow-casting in Shadow is a mereological sum, there may be some question as to *which* mereological sum is the shadow-caster. I had assumed that O_1 and O_2 were the only objects around, modulo the light source and whatever it is that the shadow is cast upon. But this can't be quite right, strictly speaking, since if O_1 and O_2 are separated by some distance, then there has got to be *something* between them—air molecules, spacetime points, etc. Whatever it is that is between O_1 and O_2 , someone might argue, is also a candidate for parthood, and hence a candidate for being part of a mereological sum. Also, since we did not assume that O_1 and O_2 are mereological simples, O_1 and O_2 each have (at least) a front half and a back half. In which case, as concerns Figure 1, there is the mereological sum of O_1 and O_2 , the mereological sum of O_1 , O_2 , and all of the spacetime points in between

O_1 and O_2 , the mereological sum of O_1 , O_2 , and only *some* of the spacetime points between O_1 and O_2 , the mereological sum of the front half of O_1 , the back half of O_2 , and every other spacetime point in between the two halves, and so on. So even if we grant that some mereological sum or other casts the shadow in Figure 1, it need not be the mereological sum of O_1 and O_2 , since there are numerous other sums, all of which are equally appropriate candidates for shadow casting (so one might argue).

I do not wish to belabor the above point too long, since I think an appeal to simplicity will rule out many of the arbitrary sums mentioned above. Moreover, since spacetime points (or air molecules or whatever it is that is between O_1 and O_2) do not usually cast shadows (e.g., none of them are opaque, for example, and they are just not the sort of things that *could* effectively block light, or cast shadows at all), then we lack motivation to think that such things play a part in shadow-casting, by being part of a mereological sum that shadow-casts, for example. Finally, the complaint launched above borders on a purely epistemic matter, not a metaphysical one: the above objector grants that a mereological sum is involved in casting a shadow, but is worried that we have no principled reason for picking out the right one (among the many viable candidates). But this just amounts to a worry that we don't (or can't) know which mereological sum (among the many) casts the shadow, even if we grant that one of them surely does. And my main interest is to convince the reader that there are indeed mereological sums, and that these sums are just identical to the parts that compose them. So an epistemological worry about our access to these sums, or how we can tell one of them is causally efficacious in a

particular scenario over another is orthogonal to my aims here. As such, resistance to Premise 5 along these lines will be ineffective.

My main motivation for launching the above argument above is to impress upon the reader that mereological sums are not merely helpful in solving a puzzle such as Shadow, but—even stronger—a proof of their existence might be concluded from such a puzzle.¹⁵

The usual reason people resist an argument such as the one outlined above is because of concerns about ontological extravagance.¹⁶ If CI were false, then claiming that the mereological sum of O_1 and O_2 casts the shadow multiplies our ontological commitments exponentially, especially since the shadow example generalizes to any of the usual cases of shadow-casting. And this is why many in the literature have shied away from mereological sums as a solution to the puzzle—it would seemingly carry with it too heavy of an ontological burden, given the ubiquity of shadows. Not so with CI, however, as I hope I have already shown in previous chapters. If composition is indeed identity, then we can appeal to mereological sums to solve puzzles such as the Shadow Puzzle without accruing any extra, unwanted ontological costs. And our conclusion, 9, of the argument above will seem appropriately benign, as it should.

Before moving on to the next puzzle, I'd like to take a moment to compare Shadow with the two cases of causation I discussed in the previous section:

¹⁵ Unless, of course, as I explained above, you accept 5a but deny 5b, in which case our discussion will take a decidedly different turn, and we will be debating other arguments presented in this thesis.

¹⁶ See Carolina Sartorio (forthcoming) "Failing to do the Impossible" in *New Waves of Philosophy of Action*, and (2006) "Disjunctive Causes" in *Journal of Philosophy* 103. See also: Roy Sorensen (2008) *Seeing Dark Things*.

collective causation and overdetermination. We have already seen how O_1 and O_2 casting the shadow shares some features of our paradigm case of collective causation (the bomb case): both O_1 and O_2 in Shadow and the men turning the key in the bomb case are objects—plural—that engage in some activity collectively. Since it is unproblematic to claim that both men can, collectively, cause the bomb to go off, it should be unproblematic to claim that both O_1 and O_2 can, collectively, cast a shadow. But as our discussion a few pages up revealed, there is more to the story. For, in some ways, our Shadow puzzle is more like a case of overdetermination than it is a case of collective causation. Here's why. In the collective causation case, the following two counterfactuals hold: if ONE would not have turned the key, and TWO would have, the bomb would not have gone off, and if TWO would not have turned the key, and TWO would have, then the bomb would not have gone off. In other words, in the bomb case, each man is necessary but not sufficient for the bomb's detonation. In Shadow, however, the following two counterfactuals hold: if O_1 would not have been there, but O_2 remained where it is, then the shadow still would have been cast, and if O_2 would not have been there, but O_1 remained, then the shadow still would have been cast. In other words, each of O_1 and O_2 is sufficient but not necessary for the shadow being cast.¹⁷

There's an interesting question here about the criterion of identity for shadows. Are shadows dependent upon their casters? If we begin with a situation as represented in Figure 1, and then remove O_1 , and leave O_2 where it is, we can all agree that the shadow will *look* as if it hasn't changed. Similarly, if we were to keep O_1 in its original position and remove O_2 , it will *look* as if we haven't done a thing, if

¹⁷ Again, this was discussed briefly above, pp. 9-10.

we keep our attention focused on the shadow. As far as the appearances of shadows go, then, it will be indiscernible whether we have both O_1 and O_2 in place, or just O_1 , or just O_2 , or an entirely different opaque object that is neither O_1 nor O_2 but similar in size, or any number of different objects appropriately shaped, sized, and arranged. But if a shadow's identity is dependent on its caster, then it will be irrelevant whether many different shadows are qualitatively indiscernible: a difference in caster would make a difference in shadow, irrespective of whether or not we can tell just by looking at the shadow.

But I need not get too mired in these issues here. The counterfactuals that show the sufficiency but non-necessity of each of O_1 and O_2 as casters of a shadow (in general) are as follows: if O_1 would not have been there, but O_2 remained where it is, then a shadow still would have been cast, and if O_2 would not have been there, but O_1 remained where it is, then a shadow still would have been cast. No parallel counterfactual is true of the paradigm case of collective causation (e.g., the bomb case), and so this is an important way in which Shadow differs from a typical case of collective causation.

In paradigm cases of overdetermination (e.g., the two shooters), the elements involved are also sufficient but not necessary for bring about a certain state of affairs (event, effect, etc.). If Man_1 and Man_2 each deliver a fatal blow to Man_3 , then the following counterfactuals are true: if Man_1 would not have fired, but Man_2 did, then Man_3 still would have died; and if Man_2 would not have fired, but Man_1 did, then Man_3 still would have died. Each of Man_1 and Man_2 , then, is sufficient but not

necessary for the death of Man₃.¹⁸ Again, for our puzzle case, Shadow: if O₁ would not have been there, but O₂ remained where it is, then a shadow still would have been cast, and if O₂ would not have been there, but O₁ remained where it is, then a shadow still would have been cast. Each of O₁ and O₂, then, is sufficient but not necessary for a shadow being cast.

It is typically the ‘sufficient-but-not-necessary’ feature that is hailed as problematic in cases of overdetermination—primarily because it violates simplicity principles. Why posit two causes when one will do? Yet if this is something that Shadow has in common with cases of overdetermination, then if it is problematic in cases of overdetermination, then it will be problematic for Shadow in the same way.¹⁹

But I fail to see why overdetermination is problematic, for reasons that I will delineate in section 4 below. Essentially, I will be arguing that if CI is true, then we get to use mereological sums with abandon. Whenever we have a case of overdetermination, we can consider not the two (or more) sufficient-but-not-

¹⁸ Notice that we could make a parallel point to the one made above about the criterion of identity of shadows, and wonder instead about to the criterion of identity for events: is the death of Man₃ when it is caused *only* by Man₁ the same event as the death of Man₃ when it is caused *only* by Man₂, or when it is caused by both Man₁ and Man₂, etc.? (Admittedly, there are some dissimilarities between the shooting case and Shadow. There are two different bullets, for example, in the shooting case. But such dissimilarities are merely due to the arbitrary details of the story. The case could easily be re-described to maintain parallel structure (e.g., two psychics who can kill just by thinking about it, etc.)) In any case, I don’t think that we need to (yet) come down on the question about the criterion of identity for events. For we could generalize by talking about the cause of a death of Man₃, in general, without worrying about whether the relevant death-events under consideration are identical or not. We can then run the relevant counterfactuals in terms of this general event, and yield the result that the men are each sufficient but not necessary, comparable to objects O₁ and O₂ in Shadow.

¹⁹ Actually, whether overdetermination is problematic or not is up for debate. Some take it as a direct violation of simplicity principles (Kim (1989), Kim (1993a, b, and c), Merricks (2001), etc.). Some take it as a direct refutation of counterfactual theories of causation, or at least proof that a revision of the counterfactual theory of causation is needed (Fine (1975), Lewis (1979), Horwich (1987), etc.). Some take it as obvious, and do not argue for it at all (see Merricks (2001)). Others, however question whether overdetermination is problematic at all. See (e.g.) Sider (2003) “What’s So Bad about Overdetermination?”

necessary causes as cancelling each other out; rather we can consider the one mereological sum of the two (or more) causes (which CI claims is simply *identical* to the two (or more) causes) as an unproblematic cause, resulting in just your usual, run-of-the-mill, ordinary case of causation. For now, however, it is only important that we see how Shadow is similar in structure to overdetermination cases, in that the relevant causal elements have the particular counterfactual features that they do.

ii. *Eclipse*²⁰

Let us move on to our second puzzle. And let us modify Shadow slightly, to create another puzzle, Eclipse. Let's imagine that the light source in the above example is actually a sun. And let us imagine that the first object is one planet, P_1 , and the second object is another planet, P_2 . We observers are on a third planet, P_3 , watching P_1 and P_2 align in front of the sun, creating an eclipse. The planets and the sun are arranged in such a way that we are the farthest from the sun, then P_2 , then P_1 , and then the sun, as represented by Figure 2 (not drawn to scale):




Figure 2

²⁰ This example modified from Sorensen (2006).

Equipped with appropriate eyewear, we stare at the celestial phenomenon. But what, exactly, are we seeing? We are seeing an eclipse, surely. But what does that mean? Are we seeing P_1 ? P_2 ? Or something else?

We are not seeing P_1 , one might argue, since P_2 blocks our view of P_1 . We can imagine that P_1 is just the appropriate size and distance from P_2 , that P_2 completely obscures our view of P_1 (from our point of view on P_3). In fact, if the sun had not been behind both P_1 and P_2 , and if it had been at the right angle to illuminate the planets, P_1 and P_2 , we on P_3 would (intuitively) only be able to see P_2 . But an eclipse is like (or simply *is*) a shadow. Indeed, the reason (total) solar eclipses are only visible from certain geographical regions on the earth is because the moon's shadow is only a few miles wide. So we can simply imagine that we have been inserted into our puzzle, Shadow, and we are now wondering, from the point of view of the shadow, what we are seeing if we are looking towards the light source. We know from our puzzle that there needs to be a light source hitting an object in order for that object to be blocking light from that light source. If this is right, then we on P_3 cannot be seeing P_2 , since P_2 has no light shining on it at all. P_2 can be no more responsible for blocking the sun than O_2 in the Shadow Puzzle was responsible for casting a shadow—neither is in a line of light to cast a shadow!

Imagine that you are in a dark room of your house.²¹ There is light outside, which you can see by looking toward the window. Suppose you have a giant cardboard cutout of Obama placed in front of your window. When you look at your window from inside your dark house, you see a silhouetted Obama-shape. Now

²¹ Thanks to Keith Simmons for this example.

imagine that there is an intruder standing between you and the cardboard Obama cutout. It just so happens that the intruder is the same shape as the Obama cutout. Even so, do you see the intruder? Intuitively, no! This is why you would be truly frightened if the intruder were to move, or talk, or come closer—you didn't know (or see!) that he was there! So if you cannot see an intruder, who is completely hidden by a cutout in silhouette, then it seems—for similar reasoning—you do not see P_2 in Eclipse either. It does not matter that P_2 is closer to you than P_1 —just as it doesn't matter that an intruder may be closer to you than the cardboard cutout that is causing a silhouette. So P_2 must not be causing the eclipse (just as the intruder is not causing the silhouette).

Yet if neither P_1 nor P_2 are causing the eclipse then what is?!

As with Shadow above, I'm going to suggest that not only can a single planet cause an eclipse (as we usually accept), but that a mereological sum of planets can cause an eclipse as well. And why not? After all, if Eclipse is just Shadow from the perspective of the cast shadow, then why should the ontological facts change just because of a change in perspective? If a mereological sum can cast a shadow, then a mereological sum can cause an eclipse, especially given that an eclipse just *is* a case of shadow-casting seen from the shadow's perspective.²²

Moreover, we could create a slippery slope from the purportedly problematic situation as represented in Figure 2 to the unproblematic situation when there is just one planet creating an eclipse. Imagine that we begin with two planets, as in Figure

²² I am obviously ignoring certain details about eclipses (and shadows), such as that light might be bending around the opaque objects that are blocking the original light source, etc. I hope such details are ultimately mere noisemakers for the broad metaphysical points I am interested in here, ones involving parts and wholes, and the relation between them, for example.

2. Yet by some mysterious cosmic event, P_1 and P_2 start inching toward each other. Eventually, P_1 and P_2 begin to fuse together, and are fused together perfectly such that there is just one round planet (call it P_4 ²³). We begin (assume) by saying that neither P_1 nor P_2 causes the eclipse; we end by saying that P_4 causes the eclipse. But given our tolerance for small variation—e.g., having P_1 or P_2 a millimeter more or less closer to each other seems not to make a difference in whether either is or is not casting a shadow—and in order to avoid arbitrariness, we either need to give up that P_4 causes a shadow, or we need to say that something in the original case does.

Pushing the point further, P_4 admittedly has a front half (the half facing the light source, say) and a back half (the half facing us on P_3). Yet our reasoning that generated the Eclipse Puzzle from Figure 2 could just repeat itself at the level of the halves of P_4 : it seems that we do not see the front half, since it is blocked by the back half. But we do not see the back half either, because no light hits it, and it seems that it is occluded by the shadow of the front half. So then it seems the situations are parallel: if we have a puzzle involving the two planets, we have a puzzle involving just one (if we admit that there is a front half and a back half of the one planet).

But we *don't* have a puzzle involving just one. We understand what it is to watch an eclipse in the case of just one planet blocking a light source. We are seeing the whole planet block light, even if we admit that the whole planet is made up of a front half and a back half. And if we understand how one (whole) can cause

²³ I need not assume (or speculate or come down on) whether P_4 is identical to either P_1 or P_2 , or both, or what, for my purposes here.

an eclipse, then we understand how a mereological sum could cause an eclipse. If someone objects that there is an important distinction between halves that make up (connected) wholes and unattached or unconnected parts that make up a scattered whole, then I will repeat my response that I gave in the previous section. Attachment or connectivity is not going to make the difference between objects that can collectively cause an eclipse and objects that can't, any more than attachment or connectivity is going to make a difference for any activities achieved collectively.

And let us not forget the unproblematic case of collective causation. If we can make sense of two unattached, unconnected men collectively causing a bomb to detonate (by turning two keys simultaneously), then we can make sense of two unattached, unconnected planets, P_1 and P_2 , collectively causing an eclipse.

The solution here, of course, is to embrace CI and to maintain that in the original eclipse case (represented by Figure 2) is one where a mereological sum—*viz.*, the mereological sum of P_1 and P_2 —is causing the eclipse. As with the Shadow Puzzle, many will not want to resort to mereological sums here because of the heavy ontological costs that such entities presumably bring. However, with CI as I've defended her here, there are no such worries, and so we can appeal to mereological sums with impunity to solve the Eclipse Puzzle, as well as the Shadow Puzzle.

Now perhaps one might argue that there is a difference between the cause of an eclipse and what we see when we are looking at an eclipse. Put another way, do we see the cause of the eclipse when we are looking at the eclipse? Or is what we see when we are observing an eclipse something distinct from whatever is causing the eclipse?

Here is why you might think that the *cause of an eclipse* and *what we see when we are looking at an eclipse* are different: the cause of the eclipse is just whatever object(s) cast a shadow, whereas what we see when we are watching an eclipse is just whatever objects we would see if lighting conditions were normal, only in the eclipse case these objects just happen to be lit from behind. The cause of the shadow of the planets, as we learned from the Shadow Puzzle (assume), is just the mereological sum of P_1 and P_2 . But what we see when we look in the direction of the light source, and see the eclipse, is just whatever object is in our line of sight and would be seen if lighting conditions were normal—and that's just P_2 ! If lighting conditions were normal, and we (on P_3) were looking at P_1 and P_2 in alignment, then we would presumably only see P_2 , since P_1 is perfectly occluded by P_2 . And why should the metaphysical facts about what we see change if the only difference has been the lighting? Whether the lighting is from above, in front, or behind P_2 , the facts about what objects are placed in what order in our line of vision do not change: P_2 is closest to us, and P_1 is perfectly hidden behind it. So, this line of argument continues, what we see when we see an eclipse is just what we would see under normal lighting conditions. It's just that these objects just happen to be lit from behind (which doesn't affect what we see). So, what we see when we see an eclipse (such as the one represented by Figure 2) is not P_1 , nor the mereological sum of P_1 and P_2 , rather it is just P_2 .

This line of reasoning is misguided for (at least) two reasons. First, lighting conditions clearly *do* make a difference to what we see.²⁴ When the lights are on,

²⁴ I am assuming a “non-epistemic” kind of seeing, or a “raw seeing”, not an “epistemic seeing”, or “seeing that.” See Dretske (1979), and below for a brief discussion.

and conditions are normal, objects are lit from above or in front, and we see whatever is in our line of vision, in order. When the lights go out, we see nothing, not even something that may be right in front of our face. If lights being on or off can make a difference to what we see, then perhaps where the light is coming from can make a difference as well. And, sure, the metaphysical facts have not changed. If there are two apples in front of you, one lined up directly behind the other, so that the first perfectly occludes the second given your line of vision, then the one apple is in front of the second, no matter what the lighting conditions. Of course, this doesn't mean that you always see what's in front of what. If the lights are out, for example, you see very little (or nothing at all). If the lights are low, you may see very little, perhaps you can make out shapes and outlines. If the lighting is from behind the objects in front of you, then you will see objects in silhouette. And seeing objects in silhouette is just to see whatever is causing a shadow from the shadow's perspective. Like the Obama cutout example, if something is silhouetted against a lit window, while you are looking on from a dark interior, you will not see anything between you and the silhouetted object, even if there is such a thing (e.g., an intruder). Indeed, if an intruder wanted to hide from you, and you are sitting in your dark house. Looking towards the window, which is lit from outside, an effective place for the intruder to take cover would be to place himself between you and the Obama cutout that is silhouetted against the window—i.e., to place himself in *front* of the cutout!²⁵ So lighting conditions *do* make a difference to what we see.

²⁵ Sorenson makes this point (Sorensen (2008)).

Second, even when lighting conditions are normal we never see *all* of an object. If Adam is looking at an apple, for example, Adam will not see the back of the apple, or the insides, or the top or the bottom. Once we admit this much, however, there seems to be even less reason to think that we (on P_3) are seeing *just* P_2 in Eclipse. For once we admit that that we need not see all of the parts of an object to see an object, then it won't matter that P_1 is completely occluded by P_2 —the back of the apple is completely occluded by the front part of the apple! Yet we still say that Adam sees the apple, even if the back is occluded by the front. Similarly, someone looking at a (double) eclipse can see more than just (e.g.) the planet closest to them, even if such a planet is completely occluding another.²⁶ So it is not the case that we see only P_2 in Eclipse. And there is furthermore little reason to think that *the cause of the eclipse* (e.g., the mereological sum of P_1 and P_2) is distinct from *what we are seeing* when we are looking at the eclipse. If the sum of P_1 and P_2 are indeed causing the eclipse, then the sum of P_1 and P_2 must be what we are seeing when we are seeing the eclipse.²⁷

Of course, similar to cast shadows, and causal events, there is an interesting question about the identity conditions of eclipses: are eclipses dependent on the object(s) that make them? Imagine P_1 and P_2 are aligned as they are in Figure 2 and we (on P_3) are looking at the eclipse that they make as they are lit from behind. Now imagine P_1 begins to shrink down until it is the size of a pea. Given the set-up, however, we on P_3 would be none the wiser, since the eclipse would be seemingly

²⁶ This move is foreshadowing my take on another puzzle, Perception, which is discussed below, section iv.

²⁷ Again, where 'seeing' here is non-epistemic.

unchanged. What we see is a big, round dark circle, outlined in bright light, and this doesn't change if P_2 remains as it is. However, now imagine that P_1 re-expands back to its original size, and then P_2 begins to shrink. Again, from our perspective on P_3 , nothing seems to have changed. Of course, if eclipses are dependent on the objects that make them, then our opinion that nothing has changed means little. Lots of things have changed, and we may have been experiencing a grand succession of multiple (qualitatively identical) eclipses, depending on what the identity criteria for these things are. However, as with cast shadows and causal overdetermination, we need not settle on this matter here.²⁸

Moreover, we can deflect this issue and talk about the objects that cause *an* eclipse, in general, and generate the relevant counterfactuals that are true of the P_1 and P_2 . Similar to Shadow, and similar to cases of causal overdetermination, the following counterfactuals are true of P_1 and P_2 in the eclipse case: if P_1 would have not been there, but P_2 would have, we still would have seen (there still would have been) an eclipse; if P_2 would not have been there, but P_1 would have, there still would have been an eclipse. So each of P_1 and P_2 is sufficient but not necessary for us there being an eclipse.

²⁸ And, similarly, there is a question about the criterion of identity for silhouettes. Is the silhouette that is caused by the Obama cutout in the window the same silhouette if an (Obama-shaped) intruder stands between you and the cutout? If the cutout were to suddenly shrink, you would be none the wiser, since your experience is still of a Obama-shaped silhouette. But is this the same silhouette as the one caused by (e.g.) just the cutout, or just the intruder, or both the intruder and the cutout, or something else entirely that is also Obama-shaped. Again, I will leave these questions aside. Also, above, I have made the parallel between seeing a silhouette and seeing an eclipse, without arguing for their comparison. There may be another interesting issue here as well—e.g., whether there is any genuine distinction between seeing a silhouette and seeing an eclipse (i.e., is seeing an eclipse *just* seeing a planet in silhouette?). Again, I do not have time for these issues here, even though they may be relevant and intriguing.

We still have yet to see whether cases like this one, or cases of overdetermination in general are indeed problematic or not (which I will discuss in some detail below). But for now it is enough if we grasp the important similarities between the elements involved in Eclipse, Shadow, and the elements involved in causal over-determinism.

iii. Prevention

Let us consider a third case. Imagine that a ball is being thrown in the direction of a (currently) intact window. There is a catcher in between the window and the ball, however, who catches the ball. Yet even if the catcher would not have caught the ball, there was a wall behind him that would have stopped the ball, which is represented by Figure 3:




Figure 3

If the catcher would not have been there, the wall would have prevented the ball from hitting the window. On the other hand, if the wall would not have been there, it would not have mattered, since the catcher caught the ball, and so there was nothing for the wall to stop in any case. So it seems that neither the catcher nor

the wall prevented the window from being shattered! But surely something prevented the window from being shattered: a ball was hurled in its direction, and yet the window remains intact! So what, exactly, prevented the window from shattering?!

We can see immediately how this case is analogous to Shadow and Eclipse: there is an event that cries out for explanation (the window being prevented from shattering, the casting of a shadow, the eclipse, etc.), there is the usual, unproblematic causal element (a ball being thrown in a particular direction, a light source beaming in a particular direction, etc.), and there is the pair of objects that gives us pause (the catcher and the wall, the two occluding objects, the two planets, etc.). The reason the pair of objects gives us pause in each case is because we have failed to consider that two (or more) objects can collectively engage in an activity, including casting shadows, causing eclipses, and preventing windows from shattering.

So clearly I am going to be suggesting a parallel solution in each case. The trouble arises only when we rule out the possibility that mereological sums can be causal elements, and that the parts composing the sum can engage in collective causation—the kind of collective causation outlined at the beginning of this chapter, that we are well familiar with and find (mostly) unproblematic.

To massage our intuitions, I will appeal to the strategy employed in the first two puzzles. Let us imagine the purportedly unproblematic case of just one object (e.g., a catcher) preventing a window from shattering. If we were to remove the wall from Figure 3, then we seemingly no longer have a puzzle. This is because we understand what it is to have a single object—a catcher—prevent a window from

shattering. But the catcher is an object composed of parts. In particular, the catcher has a front half and a back half; indeed, his glove and hand each have a front half and a back half, etc. So if we think that it is unproblematic that a catcher prevents something, and we also admit that a catcher has parts, then we will be admitting of a case that is similar in structure to the our puzzle case, Prevention.

Let's grant that the catcher's mitt has a front half and a back half.²⁹ Then the front half of the mitt seemingly didn't prevent anything, since the back half was directly behind it, and (assume) if the front half wouldn't have been there, then the back half would have prevented the window from shattering. But, similarly, the back half didn't *do* anything, since no ball even hits it. So it seems that neither the front half nor the back half of the mitt prevent the ball from shattering the window. But surely something did!

Now, true, the above counterfactual *if the front half wouldn't have been there, then the back half would have prevented the window from shattering* assumes that the front half and back half of the mitt (or the various parts of the preventer, whatever they are) are each sufficient but not necessary for the prevention of the window breaking. This is stipulated so as to maintain the analog between the purportedly unproblematic case and our puzzle case, Prevention. Notice that if we do not make this assumption, then the case of a single preventer becomes a case of *collective* prevention, similar to a case of collective causation. And since, as discussed above, cases of collective causation are unproblematic, then cases of

²⁹ And let us ignore that it surely wasn't the catcher's mitt alone that prevented the window—i.e., it was probably a complex object such as the mitt, plus the hand, plus various other parts of the catcher, etc., that prevented the ball from shattering the window. The point we make in the over-simplified case of the mitt being the sole preventer can generalize to the more complex, realistic case.

collective *prevention* should be unproblematic as well.³⁰ However, it is plausible that for any ordinary, seemingly unproblematic case of prevention, where one object prevents some event, the one object admittedly has parts that are sufficient but not necessary for the prevention. Take the object minus one molecule, for example; or the catcher minus his fingernail, or the catcher minus one big toe, or the mitt minus one stitch, or the mitt minus a sliver of leather material, etc. In this way, any ordinary case of prevention can be shown to be parallel in structure with our puzzle, Prevention. So this means that either both cases are problematic, or neither of them is. And, as before, I'm going to suggest that neither of them is. What prevents us from seeing Prevention as an ordinary case is our inability to consider mereological sums as genuine objects—or our inability to consider the parts of mereological sums as some things (plural)—that can do lots of things, including preventing windows from being shattered.

Some might have agreed that in the case of Shadow, appealing to mereological sums was natural and intuitive. Indeed, what else could cast the shadow if the mereological sum didn't? And some might have thought that appealing to mereological sums in the case of Eclipse, while less obvious than in the case of Shadow, was still intuitive after some thinking about the puzzle in the right way (i.e., that Eclipse was Shadow considered from the shadow's perspective). But Prevention, some one might argue, is taking the idea of mereological sums too far!

³⁰ There are many who think that cases of causation and cases of prevention are importantly different, and this issue is usually tied up with, or seen as similar to the question about whether negative facts or events are the same as or reducible to positive facts or events. Is the fact that Pierre is not here, for example, just reducible to the positive facts of who *is* here? I don't want to commit myself one way or the other for now, and I hope the above point can be made without having to settle one way or the other.

We can have mereological sums of random objects, such as O_1 and O_2 , someone might claim. And we might be able to consider mereological sums of inanimate planets, such as P_1 and P_2 . But matters start to get fairly unintuitive if we have to consider mereological sums of persons and walls!

Moreover, someone might argue, Shadow and Eclipse are clearly cases of causation, whereas Prevention is a case of, well, prevention! And preventing something is importantly different from causing something. Or so one might argue.

As to the first point, if someone is going to object that mereological sums of inanimate objects are fine, but mereological sums of (e.g.) persons as walls are not, I will remind such an objector of all of the arguments for CI (and all of the defenses for any objections *against* CI) that I have presented throughout this thesis. Once we have embraced CI—especially as I have imagined her here in this present thesis—there should be no worry at all about embracing any mereological sums whatsoever. This includes mereological sums of animate, living organisms, such as persons, and any other random (perhaps inanimate) object you chose. If you have fully grasped the CI view being defended here, then you understand that mereological sums of objects that you have already countenanced in your ontology are literally *for free*: the mereological sums are literally identical to the objects you already accept. So there should be no resistance to the claim that there are indeed mereological sums of (e.g.) catchers and walls—especially if one has already granted me that there are mereological sums of (e.g.) opaque objects such as O_1 and O_2 , and planets such as P_1 and P_2 .

As to the second point, it may well be that prevention is something distinct from causation. As I mentioned previously in a footnote, I (unfortunately) do not have the space here to take on such a debate. But even if causation and prevention are distinct, I don't see what the issue of CI and mereological sums has to do with one and not the other. Once one admits that some objects—plural—can engage in activities collectively (e.g., carry a coffin, enter a three-legged race, meet for lunch, cast shadows, etc.), then even if preventing something is an activity distinct from causing something, it is still admittedly an *activity*. And it seems that there are very few (if any) activities that cannot be done by some things collectively. So the same reasoning that we used above to convince the reader that some things could collectively cause things (e.g., casting shadows, eclipses, etc.) would apply to cases of prevention. And this would be true even if prevention and causation are distinct activities.

So, as with Shadow and Eclipse, Prevention is seemingly problematic, until you accept CI, and allow that mereological sums (e.g., the mereological sum of the catcher and the wall) can be preventers (e.g., preventing a window from being shattered). So, yet again, we see how an appeal to mereological sums—and, indirectly, an appeal to CI—can solve (yet another) philosophical problem.

*iv. Perception*³¹

Let us look at one more puzzle in detail that might—at first—seem quite a bit different than the puzzles discussed previously. Imagine that Adam is looking at an apple. We can suppose that conditions are normal: there is good lighting, Adam is of

³¹ This puzzle is borrowed and modified from Neta (2007).

sound mind, and that his eyes are working properly, etc. We can represent such a situation by Figure 4:



Adam



Apple

Figure 4

Yet, intuitively, one does not have to see all of an object in order for us to say truly that one sees an object. Adam, for example, does not need to see the back of the apple, the inside of the apple, all around the right and left sides, or the very top or bottom of the apple, for us to say truly “Adam sees an apple.” Indeed, Adam only sees the facing surface of the apple in front of him. Following Ram Neta (2007), let us then demarcate the facing surface of the apple that Adam sees in its entirety, from the rest of the apple which Adam, admittedly, does not see. We can even ‘pull apart’ the bits of the apple that Adam does see (let’s call this ‘Facing Surface’) from the bits that he does not (let’s call this ‘the Rest’) to exaggerate the point, which we can illustrate by Figure 5:



Figure 5

Some might argue that when we pull Facing Surface away from the Rest, so that Facing Surface occludes the Rest by being some distance in front of it, we would *not* be able to truly say “Adam sees an apple.”³²

Assuming that the apple survives the removal of Facing Surface, one reason some might claim that we would not be able to truly say “Adam sees the apple” is because, even though the apple survives, the surviving apple is the object *behind* Facing Surface. Adam doesn’t see the apple (according to this line of reasoning) because *the apple* is occluded by an opaque object—the slivered-off bit of what used to be the facing surface of the apple, Facing Surface. Indeed, if we were to remove the rest of the apple completely, Adam would be none the wiser; for what

³² According to one line of reasoning, this is because there is no apple for Adam to see. Once we have removed Facing Surface from the Rest, one might claim that the object(s) that remain are no longer sufficiently intact to qualify as *an apple*. So Adam does not see an apple in Figure 5 because there is no apple for him to see. Let us assume, however, that the removal of Facing Surface is not enough to destroy the apple. We often tolerate the removal of small bits of an object, without thinking that the object has thereby been destroyed (e.g., you think that you still exist even though small bits of you have fallen off as you read this; you can survive the clipping of your fingernails; a mountain can survive the loss of sediment by erosion, etc.). So let us just stipulate that our usual tolerance for the removal of parts is in play here, and applies to the bit we’ve been calling ‘Facing Surface.’ (Of course, my own gloss on this tolerance involves thinking of objects as 5-dimensional, trans-world and trans-spatio-temporal mereological sums that don’t, strictly speaking, ‘lose’ any parts whatsoever. But let’s leave this issue aside for now; I addressed this extensively in chapter 4.)

we've been calling 'the Rest' is entirely inert as far as Adam's perceptual experience goes. So, this line of argument continues, Adam is only seeing what we've called "Facing Surface"³³—but he is *not* seeing the apple—in Figure 5, while Adam *is* seeing the apple in Figure 4.

But once we grant these two points—that Adam sees an apple in Figure 4, but does not see an apple in Figure 5—then a puzzle emerges. For, following Neta, we can imagine that we take an instance where Adam does not see the apple, as illustrated by Figure 5, and move The Rest closer and closer to Facing Surface. At some point, once The Rest is touching Facing Surface, it seems that we should be inclined to say truly that Adam sees an apple. After all, we will eventually get to a case represented by Figure 4, which we've already admitted is a situation where we can say truly "Adam sees an apple."

In other words, we can (as we have previously) run a quick step-wise argument from a situation in which we think we cannot say truly "Adam sees an apple", e.g. Figure 5, to a situation in which we can say truly "Adam sees an apple", e.g. Figure 4. We keep bringing the Rest closer and closer to Facing Surface until we are right back at Figure 4. At some point, we will have to go from a situation in which Adam doesn't see the apple, to one in which he does (as admitted at the outset by our intuitions involving the ordinary case of Figure 4). But, as with all Sorites arguments, we admit of tolerance for slight variation: it shouldn't matter whether the Rest is just a smidge closer or farther from Facing Surface. It seems

³³ As Neta points out, it is a bit misleading at this point to call the object in Figure 2 (that is in Adam's direct line of vision) "Facing Surface" since it is no longer a surface of anything (except itself). But this shouldn't distract us from the puzzle at hand.

odd to say that one tiny sliver of distance between Facing Surface and the Rest will make the difference between Adam seeing an apple and not seeing an apple. Yet that is what it seems we are being forced to do if we admit that Adam sees the apple in Figure 4, but not in Figure 5.

Now perhaps one might think that once The Rest has been separated from Facing Surface, then no amount of pushing these two together will result in our original situation, represented by Figure 4. There is a pure unity, one might claim, to organically intact objects.³⁴ If I carefully peel an apple, for example, and then stick the skin back on, the difference between my creation to an unpeeled apple will perhaps be slight (and indeed nearly imperceptible if I am careful enough) but the skin is still importantly *unattached* to the object I've created. Not so with the untarnished apple, one might argue. But this is a fact about our physical limitations. We could certainly imagine that we have a fuse gun, such that whenever we point this gun at some objects, it will fuse them together perfectly, seamlessly, so that there is no difference contact-wise between objects which have been organically fused and those which have been fuse-gun fused. And if so, then there would be no difference contact-wise between an apple which had once been separated as in Figure 5, and then fused together with our fuse-gun, and the untouched apple in Figure 4. Metaphysically, at least, a perfectly fused object is seemingly no different that one that was never separated to begin with.³⁵

³⁴ Neta discusses this point about the *unity* of objects, p. 6.

³⁵ This isn't quite right. For if we are considering 'the objects' in question as 4-dimensional space-time worms (or lumpy, 5-dimensional objects, as I suggest in Chapter 4), then such mereological sums would include 'the history' of having once been separated or not—i.e., each would have different temporal parts. Then, metaphysically, they *would* be different. But even on such a view, such 4-dimensional (or 5-dimensional, lumpy) objects would be composed of spatial and temporal (and

Moreover, as Neta points out, if we require of the objects that we see that they have unity, then it will yield the unwelcome result that it is only *unitary* objects that we really see.³⁶ Neta claims that we can see (e.g.) pairs of tomatoes, and the Milky Way Galaxy. Indeed, we can see crowds of people and heaps of sand and scattered decks of cards and piles of trash and bodies of water and clouds and dust storms and flocks of birds and so on. As another example: if I slice an apple into eighths, and then spread the slices on the table, I might request that you bring me the apple, and you would presumably have no trouble gathering up the slices and bringing the apple to me, divided and detached though it may be.

Now, true, an apple, one might claim, is not a scattered object like galaxies, crowds of people, heaps of sand, and scattered decks of cards are. Unity is important, one might continue, only when we are seeing unitary objects, such as apples.

This sort of objection seems incorrect for (at least) two reasons. First, if we are thinking of the world at the molecular level (as physicists are wont to do), or at the levels of mereological sums or gunk (as metaphysicians are wont to do), there will be no difference in principle between purported scattered objects such as galaxies and heaps of trash, and purported un-scattered objects such as apples. This is because apples, at rock bottom, are just a collection of small molecules, heaped together like a pile of trash (although composed of much smaller

modal!) parts, one of which is an object at a certain time that has its physical parts arranged in a certain way, and 'touching' in a particular way. And it is *this* object that is under discussion—an isolated temporal slice—that we can compare with another isolated temporal slice and ask ourselves if there is any significant metaphysical difference between them.

³⁶ Neta, p. 6.

components). Indeed, if we had microscopic eyes, rather than the ones we have now, the world would seem a much 'heapier' place than it does to us now, because all we would see would be small groups of molecules hovering around each other, much like we see galaxies and groups of people now. A feature such as 'unity' is supposed to be metaphysically robust, and wouldn't dissipate due to a mere change in perspective or visual ability. Since purported 'unitary' objects would seemingly lose this property on closer, microscopic inspection, there must be no such property as 'unity.'

Second, 'seeing' does not seem to be an ambiguous activity, whose necessary and sufficient conditions change depending on the object of sight. Seeing an apple, seeing a galaxy, and seeing a group of people or a pile of trash, in other words, are all involve the same activity—*seeing*. But it wouldn't be if unity is required in some cases (e.g., seeing an apple) but not in others (e.g., seeing a galaxy, etc.).

Not to mention that this notion of 'unity' is somewhat obscure in any case, as Neta rightly points out. A principle of seeing that required the objects of perception to have 'unity' would have many more problems than the few cases it may account for.

One might think that there is an important difference between seeing an apple as an apple, and seeing an apple, period. Indeed, there is often a distinction in the literature made between non-epistemic seeing and epistemic seeing.³⁷ One might see an apple, for example, without seeing *that* it is an apple. Or one might see an apple, without seeing the apple *as* an apple. *Seeing that* is epistemic, whereas just plain *seeing* is not. And *seeing as* (at least) invokes an idea or a concept of

³⁷ See Dretske (1979).

something, whereas just plain *seeing* does not. You could see a cloud as an elephant, but you wouldn't be able to do so unless you already possessed the concept 'elephant.' So, one might think that one can't see an apple as an apple unless one already has a concept 'apple.' And this concept, one might argue, contains within it the notion of unity. And that is why apples need to be unified in order to see them (so one might argue).

But I am not concerned here with *seeing as*, or *seeing that*, or any kind of epistemic seeing. I am only concerned with the ordinary, *non-epistemic* notion of just plain seeing.³⁸ And unity does not seem to be a requirement of this ordinary, non-conceptual notion of just plain seeing.

So the puzzle before us is that while we all think that in Figure 4, we can say truly "Adam sees an apple," we don't think that we can say the same thing truly in Figure 5. But there does not seem to be any significant metaphysical difference between the two cases. It isn't a matter of attachment or contact (or detachment or non-contact); attachment or contact is neither necessary nor sufficient for objecthood. As said above, we can see galaxies and heaps of sand, both of which are objects whose parts are not attached. And we could fuse any objects we please—say, a cell phone, an apple, and a desk—but this fusion wouldn't thereby create a *new* object (if the mereological sum of the cell phone, the apple, and the desk is an object—which I of course think it is—then it is an object prior to the fusing. So fusion (or contact or attachment) is neither sufficient nor necessary for objecthood.

Moreover, as mentioned above, let us not forget principles of basic physics. The notions of 'contact' and 'fused-together' are folk notions, which we all have been

³⁸ As Neta (2007) himself seems to be doing. See also Dretske (1979).

taught are inappropriate when dealing with the microscopic or the molecular. Atoms and molecules—and if we dare go so far, mereological simples or gunk—are not the sorts of things that ‘touch’ or are attached in any accurate sense.

Neta calls the underlying principle at work here the **No Difference Principle**: it can make no difference to what someone sees that there just happens to be an occluded object behind, touching, or attached to the facing surface that s/he does see.³⁹ If we admit that Adam only sees Facing Surface in Figure 4 and Figure 5, and we admit that the Rest is causally inert, and seems to make no difference in what Adam sees, then according to the No Difference Principle, Adam does not see an apple in Figure 4 (or else he *does* see an apple in Figure 5).

I take it that the real point is that our verdict in each case—Figure 4 or Figure 5—must be symmetrical, given the No Difference Principle. Perhaps there’s a Moorean argument to be made here: since more people are more certain that Adam is *not* seeing an apple in Figure 5, than they are that Adam is seeing an apple in Figure 4, then given the No Difference Principle, Adam must not be seeing an apple in Figure 4.

And, of course, what goes for Adam and his apple goes for all of us everywhere. If we can’t explain why it is that Adam sees an apple in Figure 4, but not in Figure 5—if we are unable to split the difference metaphysically between what Adam sees in Figure 4 and what he sees (or doesn’t see) in Figure 5—then we may be forced to claim that Adam doesn’t see an apple in Figure 4 (or else we may be forced to say that he does see an apple in Figure 5, which is equally unintuitive). Then, generalizing, we will have to claim that none of us ever see the ordinary

³⁹ Neta p. 9; I’ve modified and generalized Neta’s definition for my purposes here.

objects that we take ourselves to see all of the time. And this would be quite devastating to our ordinary intuitions about what we do and do not see, generally.

Neta puts the argument as follows (I've modified it to suit my presentation):

1. Adam sees the same Facing Surface in Figure 4 and Figure 5.
2. There is no difference between what Adam sees in Figure 4 and Figure 5.
3. If Adam sees Facing Surface in figure 4, then he also sees an apple in Figure 4 only if the presence of the occluded part of the apple behind (touching, attached to) Facing Surface makes a difference to what she sees.
4. Therefore, Adam sees an apple in Figure 4 only if the presence of the occluded part of the apple ("the Rest") behind (touching, attached to) Facing Surface makes a difference to what she sees.

No Difference Principle: If there happens to be an occluded object behind (touching, attached to) the facing surface that Adam sees, then it makes no difference to what Adam sees.

Adam does not see an apple in Figure 4 (from 4 and No Diff Principle).

As I mentioned before, Neta uses the above argument to argue for a subtle and interesting point about the contextual features of perceptual verbs such as 'sees'. He thinks that there is a kind of ambiguity or contextual shift in premises 1-4 that explains how we can go from true premises to a false conclusion. I am not interested in investigating Neta's claims here. Instead, I would like to pursue an alternative solution to the perception puzzle, one that relates to Composition as Identity, as I've defended the view in this thesis.

In particular, I would like to draw out the parallel between Perception and the previous puzzle cases, Shadow, Eclipse, and Prevention. In each case, it seems we were given two symmetrical options: either each of the two relevant situations (e.g., Figure 4 and 5, or Figure 3 and a case where just the catcher prevents a window from shattering, or Figure 2 and a case where just one planet causes an eclipse, etc.) are problematic, or else both are *unproblematic*. The typical move is to claim

that the ordinary case is problematic. But, as we have seen, this is because many have underappreciated the value of turning to mereological sums to solve these kinds of puzzles. And, of course, if CI is true, then appealing to mereological sums will be ontologically free of charge.

This translates into a modus tollens of the above argument. Instead of concluding that Adam does not see an apple in Figure 4, because of our admission that Adam does not see an apple in Figure 5, I would like to maintain that, rather, Adam *does* see an apple in Figure 4, because he sees an apple in Figure 5 as well. The reason I think that we resist this move is because we fail to consider ordinary objects as mereological sums, and we fail to consider that mereological sums are ubiquitous (and the universality holds). And we fail to do these things, because we have failed to recognize that CI is true.

In regards to the argument above, this move amounts to a denial of the No Difference Principle. As explained above in (e.g.) the case of the Shadow Puzzle, I want to claim that neither O_1 nor O_2 individually cast the shadow. Instead, I want to claim that the mereological sum of O_1 and O_2 cast the shadow. And this conclusion is drawn in light of the reasoning that seemingly concludes that both O_1 and O_2 are *each* causally inert. I claim that running the counterfactuals individually on O_1 and O_2 to show that they are each causally irrelevant has no bearing on the claim that O_1 and O_2 *collectively* cast the shadow—i.e., the mereological sum of O_1 and O_2 cast the shadow. Similarly, in Figure 4 and 5, I want to maintain that Adam *does* see an apple, even if we think we can run counterfactuals on Facing Surface and the Rest to show that they are (individually) causally irrelevant, or rather, that each seemingly

does not make a difference to what Adam sees. So, I want to claim that even though we may run a counterfactual on the Rest—i.e., if the Rest would have been removed (in either Figure 4 or Figure 5), then this would make no difference to Adam (i.e., his perceptual experience would be qualitatively identical in a situation in which the Rest was present, and a situation in which it was not)—this does *not* preclude the mereological sum of Facing Surface and the Rest from being an object that Adam sees.

So it is *not* the case that if there happens to be an occluded object behind (touching, attached to) the facing surface that Adam sees, then it makes no difference to what Adam sees—i.e., the No Difference Principle is false. This is because the counterfactuals that we run to see whether something *makes a difference to what Adam sees*, is similar to the counterfactuals that we run to show (e.g.) that each of O_1 and O_2 are causally irrelevant in the Shadow case. But just because (e.g.) O_1 and O_2 might be causally irrelevant in Shadow, this does not mean that the mereological sum of O_1 and O_2 is causally irrelevant. Likewise, just because the Rest might be causally irrelevant, this does not mean that the mereological sum of the Rest and (e.g.) Facing Surface is causally irrelevant. And, so, Adam can see the mereological sum of the Rest and Facing Surface, even if taken individually, the Rest makes no difference to what Adam sees.

This result may seem counterintuitive. It amounts to claiming that we see objects—e.g., apples—even when they are separated or detached, as in Figure 5. But in the background, recall, is my lump theory of objects. Ordinary objects, on my view, just *are* mereological sums of spatiotemporal and modal parts. Now, it may be

that we have certain ideas of what can happen to an apple without compromising its survival. For example, we may think that an apple can survive being cut into eighths (e.g., as before if I cut the apple into eighths and put it on the table, you would not be confused if I were to then request that you bring me the apple: you would presumably gather up the slices and bring me the apple, sliced though it may be). But we may not believe that an apple survives being thrown into a blender (for then it is applesauce, e.g.). But on my lump theory of objects, this is no problem, since whenever we are seeing ordinary objects, we are seeing vast mereological sums of spatiotemporal and modal parts. As such, we are seeing many cases of overlap—in particular, many cases of world-overlap. Some apple molecules, for example, and the apple itself, are each lumpy, trans-world mereological sums that overlap in this world (perhaps for their entire spatiotemporal career here in this world, or perhaps not). Similarly, the (trans-world) mereological sum of Facing Surface and the Rest overlaps with the (trans-world) apple. Now just *where* it overlaps, and which spatiotemporal, this-worldly parts it shares is perhaps difficult to say. Perhaps the (trans-world) apple overlaps the trans-world mereological sum of Facing Surface and the Rest in Figure 5, and so perhaps apples (in general) can survive having their facing surfaces and the remainder separated for a bit. If so, then if someone sees such an object (such as Adam in Figure 5), then someone *is* seeing an apple. If not, then this isn't because there is something puzzling about what's going on in Figure 5, rather, it's just that apples don't survive being separated from their facing surfaces.

Yet whether one accepts this lump theory of objects is incidental. The central point is that once we allow that we can see mereological sums, then Perception no longer seems problematic. Because we can grant that Adam can see mereological sums of things (e.g. Facing Surface and the Rest), even if we seemingly have reason to think that each of the parts of this sum is causally irrelevant to what he does see. Of course, this may lead one to wonder whether we see *too many* things, all of the time. Aren't mereological sums ubiquitous, and hence, indistinctive?

If CI is true, and things are pretty much as we think they are⁴⁰, then, yes, we see mereological sums all of the time. Suppose (what is likely) that there is a mereological sum of my running shoe and the Eiffel Tower—call this Shoe+. And let us suppose that we are in a permissible context where seeing my running shoe is sufficient for saying truly, “Meg sees Shoe+” (much like seeing Facing Surface might be sufficient for saying truly “Adam sees the apple”). Then when I see my shoe, I also see Shoe+. And so on for any of the other myriad of mereological sums that I might be seeing when I look at my shoe (e.g., the mereological sum of my running shoe and this dissertation, the mereological sum of my running shoe and the taco I ate for dinner last night, the mereological sum of my running shoe and my pink motor scooter, etc.). But this is no problem at all, for a couple of reasons.

First and foremost, recall that I am only concerned here with non-epistemic seeing. So we may see a bunch of things all of the time, but this need not generate beliefs about all of the things that we are seeing. That is, if we are (non-epistemically) seeing a myriad of mereological sums, we need not see that we are

⁴⁰ That is, the Existence Assumption and the Parthood Assumption are both true, etc.

seeing all of these sums, nor need we see these sums as the sums that they are. Second, seeing these sums does not require that we can distinguish one from the other—i.e., it is irrelevant that we may be unable to tell whether we are seeing (e.g.) an apple, or the mereological sum of the apple and a dust mite. Indeed, most likely, if we are seeing one, we are also seeing the other. But it is an epistemic matter, as well as a bit of a practical worry, that we need to be able to tell whether we are seeing one (e.g., an apple) or the other (e.g., the mereological sum of the apple and a dust mite). So the ubiquity of mereological sums is not a problem.

Even so, someone might argue that there is yet another worry. In particular, someone might claim that there is a striking *disanalogy* between Perception and some of the previous puzzles I've discussed.⁴¹ After all, in Shadow, for example, there were clearly four primary elements that generated the puzzle: the light source, the pair of opaque objects, O_1 and O_2 , and the shadow. Similarly, for the Eclipse Puzzle: there was the light source, P_1 and P_2 , and us on P_3 . Finally, for the case of Prevention: there is the ball, the catcher and the wall, and the (unshattered) window. Yet in Perception, there seem to be only three key elements to generate the puzzle: an observer, the facing surface of the object involved, and the occluded bits of the object.

I have two points to make in response to such an objection. First, recall all of our discussion about counting in Chapter 2, and how it relates to CI. Saying that there are four elements to the puzzle assumes a method of counting that I reject, and is strictly speaking (on my view) incorrect. It also rejects the point I've been impressing upon the reader through this entire chapter. One of the reasons we find

⁴¹ Thanks to Dave Ripley for discussion here.

puzzles such as Shadow, Eclipse, Prevention, etc., so compelling is because we have failed to take seriously the idea that mereological sums are objects—objects that can cast shadows, cause eclipses, prevent windows from being shattered, etc. Once we grant that it is the mereological sum of O_1 and O_2 that cast a shadow in Shadow, for example, then we no longer see a puzzle at all. There is simply one object—a mereological sum, made of (at least) two parts—that is casting a shadow. Still, I do understand the spirit of the objection, even if it is not strictly speaking correct, given my diagnosis of mereological sums as active participants in each case, and given my notion of counting. But even the spirit of the objection can be shown to be misguided, if we reinterpret Perception in a slightly different way.

Instead of seeing Perception as an ordinary case of ‘seeing an apple’, let us consider it instead as a case of ‘prevention from seeing.’ Imagine that Adam is looking straight ahead, as he is represented as doing in Figure 4, and that there is a cat perfectly aligned behind the apple. We can represent this situation by Figure 6 (not drawn to scale):




Figure 6

We might say that the apple is preventing Adam from seeing the cat, much like a catcher might prevent a ball from shattering a window. If the apple were not there, Adam would see the cat. But then we could remove the facing surface of the apple away from the rest of the apple, as we did in Figure 5, and then we would have a case of prevention (a case where Adam is prevented from seeing the cat), much like we did in our puzzle case, Prevention. We could represent such a situation by Figure 7:




Figure 7

Figure 7 parallels the structure of the puzzle cases Shadow, Eclipse, and Prevention. We have four elements involved: Adam, the two parts of the apple (Facing Surface and the Rest), and the cat. The apple (both Facing Surface and the Rest) seems to be the (one!) thing that is preventing Adam from seeing the cat.

If we think of ordinary cases of seeing as situations where one is *prevented* from seeing whatever is behind the object in our line of sight, then we can mirror the structure of the previous puzzles. In cases where there is not a further object to be seen—e.g., the cat in Figure 7—we can think of cases of seeing as cases of

prevention from seeing through the object. For even if there is a world with just Adam and an apple, if Adam is seeing the apple, then he is not seeing what is on the other side of the apple—i.e., nothing. We might say: the apple is preventing Adam from seeing the empty space beyond the apple. So even in cases where there isn't an object to be seen, we could still think of *seeing* as *prevention from seeing* something else (or nothing at all).

Of course, we might wonder whether such a substitution merited. Why should we think that cases of seeing are just cases of being prevented from seeing something else? And does it matter? Do we, for example, have to think that all cases of causation can be recast as cases of prevention in order to pursue this line? Unfortunately, I do not have the space here to argue for this particular point. Rather, this is merely a suggestion—just one line of argument that might be available as a response to an objector who claims that Perception is not parallel to Shadow, Eclipse, and Prevention. Whether such a response will be ultimately satisfactory will have to be pursued another time.

The main point of this section is simply to push our understanding of CI, and the application of mereological sums to various puzzles throughout the literature. We've seen how mereological sums might be able to help us in solving puzzles such as Shadow, Eclipse, and Prevention. Perhaps mereological sums may also be able to aid us in solving Perception as well. Of course, a complete treatment of Perception by appeal to mereological sums will require a bit more time and attention. But I hope the reader can at least get a feel for how such a response will go. Whether such a response is ultimately accepted will have to wait for another time,

when my goals are less directed at showing the breadth of application of mereological sums, and can focus more carefully on the details of this particular puzzle. For now, I hope it suffices to have shown the reader several ways an appeal to mereological sums will be beneficial—in philosophical areas such as causation, prevention, and perception.

4. Concluding Thoughts: Broader Application and Overdetermination in General

As has been suggested above, there is a general form or template to the above four puzzle cases. In each case we deviate from the ordinary case by introducing a pair of elements, which are seemingly doing the work that one object usually accomplishes. In Shadow, for example, there were clearly four primary elements that generated the puzzle: the light source, the pair of opaque objects, O_1 and O_2 , and the shadow. Similarly, for Eclipse: there was the light source, P_1 and P_2 , and us on P_3 . And so on. Once we have such a structure available, however, we will be able to see that there is broader application of CI than I have discussed here.

Take, for example, the Exclusion Problem in philosophy of mind.⁴² Imagine that you are an Epiphenomenalist and you think that mental properties are distinct from physical properties. However, you also think that whenever you have a mental property, there is also a corresponding physical property, which the mental property is somehow dependent on. But let's imagine that the physical property causes some event, and that is sufficient for the event coming about. But then it seems that there is nothing left for the mental property to *do*; the mental event is seemingly causally

⁴² See Malcolm (1968), Kim (1989), (1993), etc.

irrelevant. To claim that there are both mental causes and physical causes for a certain (physical or mental) event, and yet to maintain that the physical causes are sufficient for the event coming about, is to posit rampant overdetermination with respect to mental causes. To avoid such rampant, inelegant overdetermination, some argue, we should simply deny that there are such mental properties, events or causes in the first place.

Of course, in defense of the dualist or Epiphenomenalist, I would suggest that she consider appealing to mereological sums as causes. Like our puzzle cases, we could maintain that mereological sums can (unproblematically) cause an event, even if one or more of the parts of the mereological sum is sufficient for bringing about the event. In the mental case, the Epiphenomenalist could claim that the mental event or property is seemingly irrelevant in the same way that (e.g.) O₂ is irrelevant in casting a shadow in Shadow, or P₂ is seemingly irrelevant in Eclipse, or the Wall is seemingly irrelevant in Prevention, or the Rest is seemingly irrelevant in Perception, etc. But just because there is a line of reasoning (a counterfactual) that seemingly shows the causal irrelevance of one part of a mereological sum, this does not then mean that the mereological sum as a whole is causally irrelevant. So, in this case, the Epiphenomenalist could claim that it is the mereological sum of the (relevant) mental event and the physical event that causes thus-and-so. In this way, appealing to mereological sums (and, indirectly, CI) could aid the Epiphenomenalist against the Exclusion Problem.⁴³

⁴³ I should note that I do not think that Epiphenomenalism is correct—nor any mind/body dualist view for that matter. But I do not think that the Exclusion Problem is an effective argument against the view, for the reasons I've delineated above.

As another example, let us consider Merricks' argument for his Eliminativist view of ordinary objects (e.g., tables and chairs), which he launches in order to avoid rampant overdetermination.⁴⁴ Merricks claims that if composite wholes existed, they would overdetermine events which are sufficiently caused by the composite's parts—e.g., the parts of a baseball cause a window to shatter, and are sufficient for the shattering, so it would be causally redundant to claim that the (whole) baseball causes the shattering as well. Since the baseball example generalizes to all cases of part/whole causation, Merricks claims that we should deny that there are any wholes. Such rampant causal redundancy, or overdetermination, in other words, should be avoided at all costs, and we would do better to deny that there are (e.g.) baseballs and running shoes than that overdetermination is ubiquitous.

My response to such a worry is three-fold. First, if CI is true, then the parts of the baseball are just *identical* to the baseball, so the overdetermination charge doesn't even get off the ground. Second, if my lump theory of objects is correct, then while the (trans-world) baseball and the (trans-world) baseball-particles are *not* identical, they are partially overlapped in this world (i.e., they share a world-chunk, and perhaps some spatiotemporal parts as well). And while it has been assumed that overdetermination of distinct casual events is problematic, it has not been shown that overdetermination with respect to overlapped objects is problematic. If Slim and Jim, who are Siamese twins joined by one hand, decide to use this hand to press a red button and detonate a bomb, I take it that we do not think that the pressing of the button was overdetermined. There was only one hand involved! It's

⁴⁴ Merricks (2001).

just that both Slim and Jim happen to share it. Likewise, if the intersection of Columbia and Franklin is treacherous because of potholes, and because of this, an accident occurs, I take it that this is also not a case of overdetermination. There is just one treacherous bit of road—where Columbia and Franklin intersect—and it is this (one!) part that caused the accident. Likewise, if the (trans-world) baseball and the (trans-world) baseball-particles happen to overlap (in (at least) one of their world-parts), then there is likewise no issue of overdetermination. There is just the (one!) world-part that causes (e.g.) the window to shatter. So, if my lump theory of objects is true, then Merricks' argument does not go through because causation by an overlapped part of something(s) is simply *not* problematic in the way that he assumes overdetermination in general is.

Third, even if we discount my lump theory of objects, and even if the baseball is not identical to its parts, it is still the case that Merricks' worry is parallel to the cases above—i.e., Shadow, Eclipse, Prevention, and Perception—and parallel to the Exclusion Problem as well. Merricks mistakenly concludes that just because some elements involved—e.g., the parts of the baseball and the baseball—are individually seemingly causally irrelevant, then both of them (taken together or collectively) are casually irrelevant. But if one embraces mereological sums, then one might be able to consider the mereological sum of the baseball and the parts of the baseball and claim that this (one!) thing is what caused the window to shatter. Again, I myself would not claim as much, since I will embrace the lump theory of objects. But this is nonetheless a move one could make once one embraces mereological sums, and embraces them as objects that can be causally efficacious. So this is the third way

an appeal to mereological sums could dodge Merricks' arguments for the Eliminativism of composite objects such as tables, chairs, and baseballs.

Finally, consider Frankfurt cases. Imagine that Rob wants to rob a bank. Todd, an evil genius wants Rob to rob the bank, too. However, since Rob is often fickle, and changes his mind at the last minute, Todd devises a plan. Todd will place a chip in Rob's head, which can be controlled remotely. If Todd thinks that Rob might hesitate and not rob the bank, then Todd will flip on the switch and make Rob rob the bank. The big day comes, and Rob robs the bank. Todd never has to flip on the switch because Rob never once hesitates to rob the bank. Here's what seems true of Rob: if he had wanted to *not* rob the bank, he wouldn't have been able to (not rob the bank). For if Rob had decided against robbing the bank, then Todd would have flipped the chip on, and the chip would have made Rob to rob the bank. This kind of example is used as a counterexample to counterfactual accounts of free will. Such an account claims that an agent, *x*, is free with respect to an action, *a*, iff *x* could have not done *a*, if *x* had wanted to. Yet in the bank robbery case, intuitively Rob did freely rob the bank—at least, most of us would hold him morally responsible if he was caught and brought to court. But he fails the left side of the bi-conditional: Rob couldn't have *not* robbed the bank, if he had wanted to (because (e.g.) Todd would have then flipped the chip on).

Given my diagnosis in the preceding cases, I will repeat the same strategy here. I will claim that just because a counterfactual is run on an individual (e.g., Rob), this does not mean that there isn't something else that might be free, or responsible for robbing the bank, even given the counterfactual analysis of free will.

For example, consider the mereological sum of Rob and the chip (and maybe even Todd, depending on how involved he is in the chip's activation component). Could this thing—this mereological sum—have done otherwise, if it had wanted to? Well, you might think, this is a particularly strange question. First, we don't tend to think of mereological sums as agents who have wants and desires. Second, and for similar reasons, we tend to think that such things are not the sorts of thing that could be *free*, and so it is strange to talk about such a thing *being able to do otherwise*.

Fair enough. Nonetheless, we can imagine that someone might pursue this line. Perhaps they might suggest that mereological sums of objects inherit their 'agent-hood' from the parts—e.g., if we consider the agent-hood of Slim, Jim, and Tim, maybe this is just the combined agent-hood of each individual. Such considerations aren't too far-fetched, as people are often inclined to attribute agent-hood to corporations, governments, states, countries, bands, teachers' unions, etc. If groups of people can be held responsible, or be attributed agent-hood, then why not the mereological sum of Rod, Todd, and an implanted chip?⁴⁵ Thus, while Rod himself may not be morally responsible for robbing the bank, perhaps the mereological sum of Rob, Todd and the chip *is*. And maybe when agents are part of a mereological sum that is morally responsible for an action, then the agents that compose the sum should be held responsible for the action. So Rob—much like many individuals who compose a corporation—would be held morally accountable for robbing the bank, even though it was strictly speaking the mereological sum (not Rob himself) who freely robbed the bank. In this way, we could save our intuitions

⁴⁵ Admittedly the chip may not add to the agent-hood, but then adding it wouldn't detract from the agent-hood we might bequeath to such a mereological sum.

that Rob was *sort of* free (he is free in virtue of being part of a mereological sum that is free). But it is not the case that Rob himself freely robs the bank, and so it does not matter that he fails the bi-conditional of the counterfactual account of free will.

I do not expect this diagnosis to be obviously right; no doubt such an account needs some fleshing out. The point, however, is that armed with a straightforward access to mereological sums (via CI), the defender of a counterfactual account of free will now has recourse to solutions to Frankfurt cases that were heretofore simply not considered, or dismissed out of hand. Now, true, there may be other reasons not to endorse a counterfactual account of free will. But Frankfurt cases *per se* may not be the counterexample some have thought, if someone embraces mereological sums, and takes a line such as the one I was suggesting above (e.g., where mereological sums of agents can have agent-hood, free will, and can percolate moral responsibility down to the sum's parts).

Whether you are ultimately convinced by any of these applications of mereological sums (and CI) to puzzle cases throughout various different topics in philosophy is neither here nor there. The underlying point is that CI allows us *carte blanche* access to mereological sums. And once we have mereological sums on board, then we will be able to recast old dilemmas in new light. This is just one remarkable benefit of adopting CI. If such applications are carried out, and if the benefits proliferate, then this will be some reason to think that CI is true (assuming, again, that the overall utility of a view is some reason to think that it is true).

Moreover, as we hopefully have seen, overdetermination should no longer be a worry *across the board*. If CI is correct, and if universality is correct (i.e., the idea

that for any two (or more) things, there is a mereological sum of those things⁴⁶), then whenever we have two (or more) competing causes, we can think of the situation as involving the (one!) mereological sum of these causes, rather than the individual causes. If the (purportedly) competing causes are subsumed under (one!) mereological sum (of which those causes are mere parts), then our overdetermination worries will thereby dissipate. We seem to begin with the assumption that one cause is unproblematic. But, like the puzzle cases above, I maintain that there is a parallel between any (purportedly) unproblematic single cause, and a case where there is more than one sufficient cause—i.e., a case of overdetermination. For a single cause is almost always composed of parts (e.g., the single object casting a shadow has a front half and a back half, the single planet causing an eclipse has a front half and a back half, the catcher has a front half and a back half, etc.), yet we ignore this in our assessment of the single cause. Likewise, we should ignore the parts of a purported case of overdetermination, and think of it instead as an unproblematic case of singular causation—it's just that the object that is causally relevant in this case is a mereological sum (of the purportedly causally problematic parts).

Again, I do not claim that all of these puzzles should be solved by appeal to mereological sums, or that the solutions that I've gestured to above will ultimately deliver a satisfactory answer. Indeed, there might be many other considerations in each case that would lean towards an alternative solution. But the point is twofold: (i) that mereological sums are no longer to be avoided as solutions because (e.g.) they

⁴⁶ Recall that I discussed and assumed this principle in Chapter 1, section 5. I also discussed it in Chapter 4.

are too ontologically costly; on the contrary, if we accept CI, we get access to mereological sums for free, and (ii) that the application of mereological sums is impressively broad—they can be appealed to in philosophical areas such as causation, prevention, perception, moral responsibility, etc. Every sub-field in philosophy that can benefit from the adoption of CI provides more theoretical evidence that CI is not only a coherent view, but gives us great reason to think that it is true.